m-L

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/843,145	04/27/2001	Salil Pradhan	30014343 US	7305

22879 7590 08/17/2006

HEWLETT PACKARD COMPANY
P O BOX 272400, 3404 E. HARMONY ROAD
INTELLECTUAL PROPERTY ADMINISTRATION
FORT COLLINS, CO 80527-2400

EXAMINER	
ELAHEE, MD S	
ART UNIT	**PAPER NUMBER**
2614	

DATE MAILED: 08/17/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

PTO-90C (Rev. 10/03)

Interview Summary	**Application No.** 09/843,145	**Applicant(s)** PRADHAN ET AL.	
	Examiner Md S. Elahee	**Art Unit** 2614	

All participants (applicant, applicant's representative, PTO personnel):

(1) *Md S. Elahee*. (3)_____.

(2) *Allan Low (Applicant's representative)*. (4)_____.

Date of Interview: *15 August 2006*.

Type: a)☒ Telephonic b)☐ Video Conference
c)☐ Personal [copy given to: 1)☐ applicant 2)☐ applicant's representative]

Exhibit shown or demonstration conducted: d)☐ Yes e)☒ No.
If Yes, brief description: _____.

Claim(s) discussed: _____.

Identification of prior art discussed: *Scheer*.

Agreement with respect to the claims f)☐ was reached. g)☒ was not reached. h)☐ N/A.

Substance of Interview including description of the general nature of what was agreed to if an agreement was reached, or any other comments: *Examiner informed the applicant's representative that a copy of the provisional application 60/263,317 has been mailed as an attachment of this interview summary in response to the applicant's representative's request*.

(A fuller description, if necessary, and a copy of the amendments which the examiner agreed would render the claims allowable, if available, must be attached. Also, where no copy of the amendments that would render the claims allowable is available, a summary thereof must be attached.)

THE FORMAL WRITTEN REPLY TO THE LAST OFFICE ACTION MUST INCLUDE THE SUBSTANCE OF THE INTERVIEW. (See MPEP Section 713.04). If a reply to the last Office action has already been filed, APPLICANT IS GIVEN A NON-EXTENDABLE PERIOD OF THE LONGER OF ONE MONTH OR THIRTY DAYS FROM THIS INTERVIEW DATE, OR THE MAILING DATE OF THIS INTERVIEW SUMMARY FORM, WHICHEVER IS LATER, TO FILE A STATEMENT OF THE SUBSTANCE OF THE INTERVIEW. See Summary of Record of Interview requirements on reverse side or on attached sheet.

FAN TSANG
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2600

Examiner Note: You must sign this form unless it is an Attachment to a signed Office action.

Examiner's signature, if required

Summary of Record of Interview Requirements

Manual of Patent Examining Procedure (MPEP), Section 713.04, Substance of Interview Must be Made of Record
A complete written statement as to the substance of any face-to-face, video conference, or telephone interview with regard to an application must be made of record in the application whether or not an agreement with the examiner was reached at the interview.

Title 37 Code of Federal Regulations (CFR) § 1.133 Interviews
Paragraph (b)
In every instance where reconsideration is requested in view of an interview with an examiner, a complete written statement of the reasons presented at the interview as warranting favorable action must be filed by the applicant. An interview does not remove the necessity for reply to Office action as specified in §§ 1.111, 1.135. (35 U.S.C. 132)

37 CFR §1.2 Business to be transacted in writing.
All business with the Patent or Trademark Office should be transacted in writing. The personal attendance of applicants or their attorneys or agents at the Patent and Trademark Office is unnecessary. The action of the Patent and Trademark Office will be based exclusively on the written record in the Office. No attention will be paid to any alleged oral promise, stipulation, or understanding in relation to which there is disagreement or doubt.

The action of the Patent and Trademark Office cannot be based exclusively on the written record in the Office if that record is itself incomplete through the failure to record the substance of interviews.

It is the responsibility of the applicant or the attorney or agent to make the substance of an interview of record in the application file, unless the examiner indicates he or she will do so. It is the examiner's responsibility to see that such a record is made and to correct material inaccuracies which bear directly on the question of patentability.

Examiners must complete an Interview Summary Form for each interview held where a matter of substance has been discussed during the interview by checking the appropriate boxes and filling in the blanks. Discussions regarding only procedural matters, directed solely to restriction requirements for which interview recordation is otherwise provided for in Section 812.01 of the Manual of Patent Examining Procedure, or pointing out typographical errors or unreadable script in Office actions or the like, are excluded from the interview recordation procedures below. Where the substance of an interview is completely recorded in an Examiners Amendment, no separate Interview Summary Record is required.

The Interview Summary Form shall be given an appropriate Paper No., placed in the right hand portion of the file, and listed on the "Contents" section of the file wrapper. In a personal interview, a duplicate of the Form is given to the applicant (or attorney or agent) at the conclusion of the interview. In the case of a telephone or video-conference interview, the copy is mailed to the applicant's correspondence address either with or prior to the next official communication. If additional correspondence from the examiner is not likely before an allowance or if other circumstances dictate, the Form should be mailed promptly after the interview rather than with the next official communication.

The Form provides for recordation of the following information:

- Application Number (Series Code and Serial Number)
- Name of applicant
- Name of examiner
- Date of interview
- Type of interview (telephonic, video-conference, or personal)
- Name of participant(s) (applicant, attorney or agent, examiner, other PTO personnel, etc.)
- An indication whether or not an exhibit was shown or a demonstration conducted
- An identification of the specific prior art discussed
- An indication whether an agreement was reached and if so, a description of the general nature of the agreement (may be by attachment of a copy of amendments or claims agreed as being allowable). Note: Agreement as to allowability is tentative and does not restrict further action by the examiner to the contrary.
- The signature of the examiner who conducted the interview (if Form is not an attachment to a signed Office action)

It is desirable that the examiner orally remind the applicant of his or her obligation to record the substance of the interview of each case. It should be noted, however, that the Interview Summary Form will not normally be considered a complete and proper recordation of the interview unless it includes, or is supplemented by the applicant or the examiner to include, all of the applicable items required below concerning the substance of the interview.

A complete and proper recordation of the substance of any interview should include at least the following applicable items:

1) A brief description of the nature of any exhibit shown or any demonstration conducted,
2) an identification of the claims discussed,
3) an identification of the specific prior art discussed,
4) an identification of the principal proposed amendments of a substantive nature discussed, unless these are already described on the Interview Summary Form completed by the Examiner,
5) a brief identification of the general thrust of the principal arguments presented to the examiner,
   (The identification of arguments need not be lengthy or elaborate. A verbatim or highly detailed description of the arguments is not required. The identification of the arguments is sufficient if the general nature or thrust of the principal arguments made to the examiner can be understood in the context of the application file. Of course, the applicant may desire to emphasize and fully describe those arguments which he or she feels were or might be persuasive to the examiner.)
6) a general indication of any other pertinent matters discussed, and
7) if appropriate, the general results or outcome of the interview unless already described in the Interview Summary Form completed by the examiner.

Examiners are expected to carefully review the applicant's record of the substance of an interview. If the record is not complete and accurate, the examiner will give the applicant an extendable one month time period to correct the record.

Examiner to Check for Accuracy

If the claims are allowable for other reasons of record, the examiner should send a letter setting forth the examiner's version of the statement attributed to him or her. If the record is complete and accurate, the examiner should place the indication, "Interview Record OK" on the paper recording the substance of the interview along with the date and the examiner's initials.

Please type a plus sign (+) inside this box → [+] 01-23-01

A/prov

PTO/SB/16 (8-00)
Approved for use through 10/31/2002 OMB 0651-0032
Patent and Trademark Office, U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PROVISIONAL APPLICATION FOR PATENT COVER SHEET

This is a request for filing a PROVISIONAL APPLICATION FOR PATENT under 37 CFR 1.53 (c).

jc961 U.S. PTO 01/22/01

INVENTOR(S)

Given Name (first and middle [if any])	Family Name or Surname	Residence (City and either State or Foreign Country)
Robert H.	Scheer	Chicago, Illinois

☐ *Additional inventors are being named on the _____ separately numbered sheets attached hereto*

TITLE OF THE INVENTION (280 characters max)

SYSTEM AND METHOD FOR PROVIDING INTEGRATED SUPPLY CHAIN MANAGEMENT

jc818 U.S. PTO 60/263317 01/22/01

CORRESPONDENCE ADDRESS

Direct all correspondence to:

☒ Customer Number [25541] → *Place Customer Number Bar Code Label here* 25541 PATENT TRADEMARK OFFICE

OR *Type Customer Number here*

☐ Firm or Individual Name					
Address					
Address					
City		State		ZIP	
Country		Telephone		Fax	

ENCLOSED APPLICATION PARTS *(check all that apply)*

☒ Specification *Number of Pages* [107] ☐ CD(s), Number [ ]

☒ Drawing(s) *Number of Sheets* [15] ☐ Other (specify) [ ]

☐ Application Data Sheet. See 37 CFR 1.76

METHOD OF PAYMENT OF FILING FEES FOR THIS PROVISIONAL APPLICATION FOR PATENT *(check one)*

☐ Applicant claims small entity status. See 37 CFR 1.27.

☐ A check or money order is enclosed to cover the filing fees

☒ The Commissioner is hereby authorized to charge filing fees or credit any overpayment to Deposit Account Number: [011,156]

FILING FEE AMOUNT ($) [$150.00]

☐ Payment by credit card. Form PTO-2038 is attached.

The invention was made by an agency of the United States Government or under a contract with an agency of the United States Government.

☒ No.

☐ Yes, the name of the U.S. Government agency and the Government contract number are: _____.

Respectfully submitted,
SIGNATURE Gary Jarosik

Date [1/22/01]

TYPED or PRINTED NAME Gary R. Jarosik

REGISTRATION NO. (*if appropriate*) [35,906]

TELEPHONE (312) 715-4000

Docket Number: [31083.05P1]

USE ONLY FOR FILING A PROVISIONAL APPLICATION FOR PATENT

This collection of information is required by 37 CFR 1.51 The information is used by the public to file (and by the PTO to process) a provisional application Confidentiality is governed by 35 U S C 122 and 37 CFR 1 14. This collection is estimated to take 8 hours to complete, including gathering, preparing, and submitting the complete provisional application to the PTO Time will vary depending upon the individual case Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U S Patent and Trademark Office, U S Department of Commerce, Washington, D C, 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS SEND TO Box Provisional Application, Assistant Commissioner for Patents, Washington, D C 20231

EL684657798US

Reviewed and revised by Robert H. Scheer on Decemer 19, 2000

SYSTEM AND METHOD FOR PROVIDING INTEGRATED SUPPLY CHAIN MANAGEMENT

BACKGROUND

This invention relates generally to supply chain management and, more particularly, relates to a system and method for providing integrated supply chain management using networked computer systems.

In the global economy of today, supply chains are commonly used to deliver goods reliably and at affordable prices. A supply chain typically involves the flow of material, information, and money between customers, suppliers, manufacturers, distributors and, possibly, financial institutions. Material flow includes, among other things, the physical product flows from suppliers to customers through the chain and reverse flows via product returns, servicing, recycling, and disposal. Information flow involves order transmission, order confirmation, and delivery status. Monetary or financial flow includes credit terms, payment schedules, payments, and consignment and title ownership arrangements. These flows cut across multiple functions and areas both within an organization and across organizations. In this regard, supply chains exist in both service, retail and manufacturing organizations, although the complexity of the supply chain may vary greatly from industry to industry and organization to organization. The coordination and integration of the material, information and financial flows within and across organizations is critical to effective supply chain management. Thus, in order for the supply chain to work for its intended purpose, all organizations involved in the supply chain must coordinate their activities with one another so that efficiency throughout the supply chain is achieved.

To coordinate activities within a supply chain, Manufacturing Resource Planning ("MRP") and Enterprise Resource Planning ("ERP") tools have been employed by organizations in an effort to gain control of the flows, plant operations, and to provide management with timely and useful reporting. Some companies use explicit supply chain management (SCM) and supply chain execution (SCE) systems which are often focused on a specific functional requirement. Generally, MRP systems have been used to translate the schedule for the production of products into time-phased net requirements for the sub-assemblies, components and raw materials, planning and procurement. ERP systems address the technology aspects of MRP such as client/server distributed architecture, RDBMS, object oriented programming etc.

In addition to MRP, ERP and SCM/SCE systems, a further supply chain management tool is described in U.S. Patent No. 6,157,915 which describes an active collaboration technology in an open architectural framework that is used to deliver information and decision support tools to various organizations. In the framework described in the '915 patent, the people across the organizations collaborate through domain task and specific active documents. The active documents contain both the necessary business information and decision support tools. Dynamic decision making is thus made possible through the delivery of active documents to appropriate parties in response to events that are triggered by business processes, the organizations themselves, or other applications. In this manner, the active documents allow organizations to exchange information, use decision tools to act on the shared information, respond to dynamic events that require decision making, and engage the proper role players in accordance with the business process.

The method and means described in the '915 patent provide the most value in a manufacturing supply chain where people from all the supply chain partners must collaborate to create the production and distribution plans (called "business scenarios" in the patent). The '915 patent focuses on user access security, workflow routing of the "active documents" (aka Lotus Notes documents) and the inclusion in those documents of links to data warehouse information sources and decision support tools which the users can utilize in defining, analyzing, modifying, and approving the business scenarios. These links can point to the partner's internal systems which the partner has authorized the appropriate people from the other supply chain partners to access.

In contrast, the present invention applies to supply chains that rely more on automated planning and execution. While the '915 business scenario represents a statement of requirements and plan for an entire production run, the advance demand notice of the present invention is used for a single business event or transaction, namely a scheduled maintenance job at the customer's end of the supply chain. Where humans do the analysis, planning, and approval in the '915 patent, intelligent software agents and other system programs accomplish those tasks in the present invention. The present invention also embraces the complete supply chain execution of the fulfillment plan associated with the advance demand notice where the '915 patent covers only the planning and definition of the business scenario (refer to Figure 9 of that patent). The present invention provides the best value in supply chains for maintenance, repair, and operating supplies (MRO) and other supply chains where supply chain management and execution is done as a result of specific business transactions between the customer and

the supply chain and the supply chain's fulfillment of the customer's requirement is not always done on a routine, repetitive, or pre-defined manner.

SUMMARY OF THE INVENTION

The supply chain management system and method of the present invention will allow companies to operate an entire supply chain on a "just in time" basis without requiring those companies to keep a full level of product safety stock on hand. The supply chain management system will achieve this by providing a system that can forecast the demand for products thereby allowing a supplier company to replenish products in a "just in time" manner. This will reduce the amount of inventory throughout the supply chain without reducing the level of service fulfillment to the customers.

The supply chain management system will achieve this goal by using intelligent agents that are distributed across the supply chain. The intelligent agents allow for the monitoring and managing of state changes in the supply chain. The supply chain management system also includes a database of forecast data useful for determining the amount of inventory needed by a particular supplier. The forecast data may be comprised of the following:

(1) Expected consumption rates based on a historical data. The expected consumption rates can be developed by tracking actual consumption rate data for customers and compiling a database of those actual consumption rates. This can be accomplished by providing a collaborative relationship between the supplier company

and the customer and continually updating the actual consumption rate data and the expected consumption rates;

(2) Deterministic demand data based on scheduled maintenance for customer equipment/facilities and an equipment knowledge database comprised of reliability data, maintenance requirements and completed maintenance job data. The deterministic demand data may also include information related to expected repair parts, contingent repair parts, ancillary supplies and necessary tools. The customer may also choose to maintain a small "just in case" inventory or to receive "just in time" delivery from the supplier company. (Note: to achieve the best results, direct shipment from the manufacturer to the customer may be used.) Again, a collaborative relationship will be required between the customer, the manufacturer and the supplier company for developing and maintaining a deterministic demand database that is populated with deterministic demand data, as described above; and

(3) Non-Deterministic demand data based on the equipment knowledge database described above. The non-deterministic demand data is formulated by comparing historical maintenance history for customer equipment to scheduled maintenance plans. The difference between the historical maintenance demand level and the preventive maintenance demand level produces the non-deterministic demand data. In contrast to the basing inventory on aggregate market forces, the non-deterministic demand data is advantageous because it is compiled at a customer-specific level which makes the data more useful and precise.

These databases are populated by data from the customer, the manufacturer, the supplier and possibly third party data compilers/publishers. For example, the customer

will supply scheduled maintenance information (this may include scheduled preventive maintenance, condition-based maintenance, planned maintenance projects, etc.), historical maintenance results and advanced order information, the manufacturer will provide equipment maintenance, replacement information and data on anticipated reliability, and the supplier will supply actual consumption rate information. The supplier company will be in direct contact with the manufacturer for fulfilling product orders and with its branch offices and other distribution points for delivering the products to consumers. The manufacturer will be connected to the supplier company and the equipment database via the Internet or a similar communications medium.

In operation, assuming that the customer has cataloged its equipment and similar inventory into the Computerized Maintenance Management System ("CMMS" and alternatively called Enterprise Asset Management System or "EAM"), the customer begins by creating a work order. The work order can result from automated monitoring of the customer equipment, preventive maintenance planning, periodic and routine maintenance schedules, planned maintenance projects or unplanned equipment failures. The work order may be stored within the CMMS system and includes information such as possible repair parts, consumables, supplies, and tools and equipment needed for the task. An intelligent agent works in connection with the CMMS system by monitoring for any entered or modified work orders. The intelligent agent may extract appropriate data from the CMMS system to create an advanced demand notice record for products that may be required for a particular maintenance task.

The supply chain management system will receive the advanced demand notice from the intelligent agent and a further team of intelligent agents will collaborate to

determine the probability that the product(s) will be needed. This determination is made, in part, by using the forecast data. The intelligent agents will also determine if the customer already has the product in house or if the product will be needed from a supplier. The advanced demand notice is updated with the probability information and forwarded to a supplier.

At the supplier, an advanced demand intelligent software agent will accept the advance demand notice and determine a fulfillment level for the product. Examples of fulfillment levels include the "Level of Service" and the "Speed and Convenience" models. "Level of Service" refers to planned purchases when the customer can designate a point in the supply chain where product should be staged and reserved for anticipated use based on a trade-off between the risk of untimely availability and the staging costs. "Speed and Convenience" refers to unplanned purchases when the customer requires immediate product availability and delivery. The fulfillment level for a particular product may be dependent upon several factors, such as the amount of lead time preceding a particular maintenance task, the inventory supply category for the product, and other factors determined by the supply chain management system.

By way of example, if the fulfillment level is the "Speed and Convenience" model, then Buy-Hold-Sell ("BHS") procedures and practices will apply. Moreover, if the probability of demand is 100% in the "Speed and Convenience" model, then the product will be shipped from the supplier to the customer. If the probability is less than 100%, the product will be reserved for this demand, but kept at its current location.

Based on a plurality of sourcing factors, the intelligent agents should determine sourcing alternatives for the advanced demand notice. The sourcing factors may include

factors, such as, the required fulfillment level for a particular customer, the inventory supply category for the product, the lead time required for the maintenance project, and the number of sourcing suppliers capable of providing the product to the customer according to the customer's time and delivery constraints. If an automated sourcing alternative cannot be generated, the intelligent agent should contact a designated customer agent to select available sourcing options on an offline basis. If the designated customer agent requires additional approval by the customer, the designated customer agent may be in charge of coordinating the approval procedure at the customer site. The designated customer agent can be contacted via telephone, e-mail or similar communication means. All off-line decisions by the designated customer agent will be entered into the customer system and synchronized with the supply chain management system, the supplier computer system, the manufacturer computer system and any other computer systems attached to the supply chain management system that track customer purchase orders.

For automatically generating sourcing options, the intelligent agent may first determine each supplier's ability to provide the product. For example, the product may be held on an available to promise ("ATP") or capable to deliver ("CTD") status. The various sourcing options are established manufacturing supply chain methods and practices. For example, if the product is being sourced from the supplier, then a fulfillment level may be provided within the advanced demand notice. The available fulfillment levels may start back ("up stream") as far as possible in the supplier network. Once the supply chain management system validates a particular sourcing option, the supply chain management system may then issue a purchase order for the product in

accordance with the Level of Service requirements provided within the advanced demand notice constraints.

A fulfillment plan will be attached to the advanced demand notice and any products retained within either the supplier's or the supply chain management system's inventory will be committed as part of the fulfillment plan. The supply chain management system will monitor the fulfillment progress and may make this information available to suppliers, manufacturers and customers. The fulfillment status may be monitored by authorized parties to allow authorized parties to track delivery of the products.

The intelligent agent will also verify that the fulfillment plan is being executed according to the Level of Service requirements provided by the customer. If the fulfillment plan is not properly executed, the intelligent agent will generate responsive and/or corrective fulfillment plans. If no corrective fulfillment plans can be generated, then the intelligent agent will contact the designated customer agent and formulate an offline sourcing solution for the fulfillment plan. The intelligent agent may also be in charge of processing any changes in the advanced demand notice constraints and generating amended fulfillment plans based on those changes. The advanced demand notice constraints may also include payment terms for the product. The payment terms may include information, such as form of payment and the timing of the payment. For example, payment may be due when the customer receives the shipment or when the maintenance task is completed. In addition to tracking the status of the fulfillment plan, the intelligent agent may also initiate billing cycles and send invoices to customers.

If the probability of demand is 100% for a particular product, the customer may use the advanced demand notice as the purchase order or issue an independent purchase order. If the advanced demand notice serves as the purchase order, the product may be automatically delivered to the customer. If the probability of demand is less than 100% or if a separate purchase order is needed, the product will remain at a delivery site in accordance with the level of service agreement. When the customer determines that the product is actually needed, a purchase order will be submitted electronically via the customer computer system in connection with the supply chain management system. Electronic submission of purchase orders allows the supply chain management system to process purchase orders on a real-time basis. The supply chain management system will then provide the product to the customer by releasing the product for pick-up from a supplier branch, shipping the product to the customer from a supplier distribution center, providing for drop shipment from the supplier to the customer, or issuing the product from on-site inventory which is managed by the supply chain management system.

If the probability of demand is less than 100%, the advanced demand notice may also indicate how long the product should be held at the final staging point before a purchase order must be received from the customer computer system. For example, if the probability of demand is only 50% and a purchase order is not received within a designated time, then the supply chain management system will assume that the product was not needed in the maintenance task and determine what should be done with the product. Depending on the contractual agreement with the customer, there may be a restocking fee for unneeded products.

When a maintenance task is complete, the intelligent agent will capture pertinent maintenance data which includes, among other things, lists and quantities of products used during the maintenance work and may also include statements of equipment condition before and/or after the maintenance work. The intelligent agent will then pass this maintenance data to the supply chain management system and the maintenance data will be added to the forecast data for the supply chain management system. The intelligent agent may also synchronize the maintenance data with the inventory supply data for allowing the supply chain management system to release products from the final staging points or reservation status.

MRO ISCM - Inventory Management Processes

The scope of inventory management as described in this document is to forecast product demand and establish base stocking levels and re-order points for the product throughout the Grainger logistics network. The processes described are presented in a generalized manner so that they can be applied to any logistics network topology, not just Grainger's network. Consequently, they are implemented using a high degree of table-driven and parameter driven software engineering techniques.

Since the actual mathematical concepts behind forecasting are not patentable, this document does not include a specification of which mathematical forecasting approach to use under various circumstances. Rather, it provides the process framework in which the mathematics will reside. It is this process framework which is sought for protection through the patent claims.

The entire emphasis on these inventory management processes is to allow the distributor to deal with probability and uncertainty in demand. In the MRO environment, demand is much more uncertain than in the manufacturing arena (where demand is based on production schedules and plans) or in the retail arena where demand is done in large quantities. In the MRO environment, most product (especially for product in the repair parts category) experiences very small inventory turns and demand, and this demand is very uncertain depending the customers' needs. It is this ability to deal with demand uncertainty that differentiates these inventory management processes from those of other kinds of businesses.

Assumptions About the Logistics Network

These processes assume that the inventory management is applied to a n-tier logistics network. This means that there are at least two levels of distribution or order fulfillment points between the supplier and the customer. Each distribution point or fulfillment point will be designated as to whether it is replenished by a higher tier point in the network or directly by the supplier (both options are allowed for a given distribution point) and whether it fulfills customer orders or replenishment orders for a lower tier point in the network (both options are allowed for a given distribution point).

The product information data base will be annotated as to restrictions placed on individual products and their allowed locations in the n-tier logistics network. For example, Hazmat items may only be able to be shipped from certain facilities and slow moving items may only be stocked in certain distribution centers.

Demand History

Demand history is maintained within the system at a unit of time granularity determined by a system parameter. Examples are history aggregated by month, by quarter, by week, or by day. The choice of time unit depends on the overall size of the database (volume of data implications on processing time) and the overall precision of the forecasting methodologies (some may only provide forecasts that are good enough at a monthly level, for example).

Demand data is kept for each distribution point and broken down as follows:

- Demand for replenishing lower tier distribution points
- Demand for fulfilling customer orders
    - Demand for fulfilling speed and convenience orders
    - Demand for fulfilling advance demand notices

Product Segmentation

Each product will be categorized across the following segmentation dimensions:

- **MOVING CATEGORY** – A product will be considered to be either Fast Moving or Slow Moving based on the level of demand for the time period. A system parameter defines the demand level threshold that separates the two categories.

- **DEMAND RATE** – The rate that product experiences demand over a period will either be considered Fixed or Variable. Demand, at this point, is the arrival of customer orders. The determination is made by computing the standard deviation of the demand rates over the periods within the history window. A system parameter defines the threshold of standard deviation as a percent of the mean that will be the dividing point between the two demand rate categories. For example, the inventory management team may consider any product with a standard deviation of 15% or less of the mean demand to be considered a fixed rate and the mean will be used as that demand rate. If demand is considered to be variable, then the historical data is run through regression analysis to determine if it can be "best fit" to one of the following stochastic distribution functions: Poisson, Exponential, or Gamma. If a company chooses to embrace other stochastic distribution functions, there is nothing in this process definition to prohibit that. The objective is to match the product demand history to some kind of stochastic distribution function which will subsequently be used in the forecasting process.

- **NUMBER PER ORDER** – The number of units of product on a single order will either be a constant (usually 1 in the MRO environment) or will be Lumpy. As was done for Demand Rate, a standard deviation is computed for the quantity of units of the target product per order. This standard deviation expressed as a percent of the mean is compared to a system parameter threshold. If the standard deviation as a percent of the mean is below the threshold, the number per order is considered to be a constant and the mean (rounded up to the next whole number) is that constant. If the percent exceeds the threshold, the number per order is Lumpy. The Lumpy category means that the number is random. As with demand rate, an attempt will be made to determine the best stochastic distribution function fit using regression analysis to describe the Lumpy number per order. Experience indicates that this is generally a Poisson distribution curve.

- **WORLD FACTORS** – The product is annotated as to whether its demand is impacted by external world factors such as the weather, the economy, competition, change in customer status (where demand comes from only a small number of customers), etc. Due to the complexity of dealing with World Factors in the forecasting process, this category will be restricted to the designation of only one World Factor (the most significant factor). If the product is not impacted by a World Factor, then its designation will be "None".

- **LEADTIME** – The product's lead-time from the supplier will be designated as either Fixed or Variable. Much of this depends on the agreements that are negotiated between the distributor and the supplier and on the supplier's capability to either Deliver-From-Stock or Deliver-To-Order at a fixed lead-time. As with the demand rate, variable lead-time history will be regressively analyzed to determine the best-fit stochastic distribution function. Experience indicates that the Poisson distribution generally works. However, there are examples where the Exponential distribution function is a better fit.

For those dimensions which have a variable alternative, the product information database will contain the designation of which stochastic distribution function is the best fit along with the applicable parameter values for that distribution function.

Taking the value combinations for these dimensions (Moving is Fast or Slow, Demand is Fixed or Variable, Number per Order is Fixed or Lumpy, World Factors is Yes or No, and Lead-time is Fixed or Variable), all the combinations are considered. There will be at least one mathematical or heuristic forecasting method applicable to each combination. Some combinations will have several possible methods. Also, any particular mathematics or heuristic forecasting method may be applicable to more than one combination. These assignments are kept in system parameter tables so that the inventory management team can easily change them. Since the particular mathematics and heuristic forecasting algorithms and methods are not the subject of the patent, they will not be discussed in this paper. Numerous books and professional articles are written on mathematical and heuristic algorithms and methods for inventory management. One that covers many of the situations found in the MRO Supply Chain domain is: *Foundations of Inventory Management* by Paul H. Zipkin (McGraw-Hill, 2000 ISBN 0-256-11379-3). In this book, Prof. Zipkin covers forecasting in end-to-end supply chains where stochastic factors are prevalent.

Inventory Management Process

The inventory management process consists of six major steps:

1. Compute and aggregate the historical demand data to be used in forecasting.
2. Forecast the combined demand.
3. Determine the critical stocking ratio that will indicate the total quantity the company can afford to hold in inventory during the forecast period.
4. Allocate the permitted inventory level from step 3 among the various distribution points in the logistics network by assigning by period the base stock level and the reorder point for each product at each distribution point. If a product is not to be allocated to a particular point, then the base stocking level for that point will be zero.
5. Determine the replenishment method that will be used for that product at the distribution point. If the use-one, replenish-one method is used, then the reorder point will be set to the base stock level minus 1. If the "(r,s)" policy is used, the re-order point, "r", will be set to whatever the chosen mathematical or heuristic algorithm determines.
6. Run the initial replenishment processes that will create orders at any point where the unreserved on-hand plus orders-in-transit quantities are less than the designated reorder point. This initial run will also determine the disposition of any product that has been reclassified as excess because the new base stock level has been reduced or set to zero and the unreserved on-hand inventory plus orders-in-transit exceeds the new base stock level.
7. Subsequent replenishments will be based on the current unreserved inventory position and the setting of the re-order points.

System parameters are used to determine the size of each forecast period. (Generally forecasting will be done in monthly periods unless the inventory management team has a particular reason to choose a different period size). Other system parameters are used to determine how many periods to include in the forecast horizon. Generally, this will be either 12 or 13 months, although there are many situations where the inventory management team will want to forecast only for a quarter or half-year.

Other system parameters determine the frequency with which history data is extracted, aggregated and loaded in the database(s) that will drive forecasting. These activities will generally be tied into the

company's data warehouse Extract-Translate-Load (ETL) schedule since the same data is used for decision support analysis as is used for forecasting.

Still other system parameters determine the frequency with which forecasting is done and which products are included in each forecasting cycle. Usually in the MRO environment, the product database embraces a large number of products and, for pragmatic purposes, the forecasting system(s) can only work on a portion of the total database during one cycle. In general, the frequency and cycle assignments should be set up so that each product is re-forecast at least annually. Many times the inventory management team will want to re-forecast semi-annually or quarterly. A few products may warrant re-forecasting monthly. This simply points out the need to have all of these scheduling and assignment related attributes be parameter driven so that the inventory management team can easily change the attribute values.

Compile Historical Demand Data

Historical demand data must be granulated across the various source dimensions (Tier Replenishment, Speed and Convenience Customer Fulfillment and ADN Fulfillment) as well as for the defined forecasting period (monthly, for example).

First, demand history must be classified as either Fast Moving or Slow Moving. Parameters are used to define the threshold demand over specified time periods for a product to be classified as either Fast or Slow Moving. These parameters can be defined differently for different distribution points in the logistics network. Also, a particular product might be classified as slow moving at lower tiers in the network hierarchy and as fast moving in the higher tiers.

Next, the demand rate (the rate at which customer orders are received) is determined. This is determined by calculating the mean and standard deviation of periodic demands of the history horizon. For example, demand might be granulized to monthly periods over a two-year history horizon. A system parameter is used to determine the threshold value for considering demand to be fixed or variable. This threshold is expressed as the percent of standard deviation to mean. For example, the inventory team may declare that if the standard deviation is 10% or less of the mean demand, then the demand rate is considered to be fixed. If the demand rate is determined to be variable, then the regression analyses are run to determine which stochastic distribution function provides the best fit for modeling the variability in demand rates (along with the applicable attribute parameters for that distribution function).

Next, the number of items per order is analyzed to determine if the order quantity is "lumpy". As was done with demand rate, the mean and standard deviations are calculated for all orders involving line items with the subject product over the historical horizon. The standard deviation is again expressed as a percentage of the mean, and a system parameter will define the threshold of this percentage to classify the number of items as constant or lumpy. If the number of items is lumpy, then the regression analyses will be run to determine the best fit and applicable parameters for the possible stochastic distribution functions to model the lumpiness.

To consider the historical effect of world factors, it is necessary for the company to have recorded the particular world factor events for each historical period. Each type of event has its own method of recording and association with demand. For example, with hurricanes, the occurrences are measured in "events". For severe summer heat, the world factor can be measured in degree-days for the region serviced by the distribution point. The parameters which define each world factor will also designate how it is to be associated with the demand. The historical data for all periods in which the world factor did not have an occurrence is compared against the periods during which such occurrences happened. The differences between these two demands represents the impact of the world factor during the historical horizon. The exact method of associating world factor affect on demand is specified by parameters. Possibly ways include "percent increase per hurricane occurrence" or a table which indicates the incremental amount of demand of product in a period based on the number of days the temperature was a certain amount above the normal. Such a table might be:

Number of Days of 10 to 15 degrees above normal	5% increase in demand per day
Number of Days of 16 to 20 degrees above normal	10% increase in demand per day
Number of Days of over 20 degrees above normal	15% increase in demand per day

Recall that the product information database has the particular world factor (if any) that applies to the individual product. This will have been established by the product management team.

Finally, the classification of lead-time as either variable or fixed must be made. As stated before, this is largely a function of the agreement negotiated between the distributor and the supplier. Some suppliers may be contractually required to deliver in 5 days, for example, or perhaps 10 days. In this case, the product information database will be so annotated and the specified lead-time recorded for the product. Where no such agreement exists, the historical data is analyzed to study the actual lead-time experienced from the supplier. The lead-times for orders placed with suppliers during the historical horizon is used and smoothed or weighted in a manner that favors the most recent orders. The mean and standard deviation of lead-times is computed. As in other categories, the threshold to determine if lead-time is constant or variable depends on a parameter stating the percent that the standard deviation is to the mean. For example, if the standard deviation is 20% or less than the mean, then the lead-times are considered constant. If the lead-times are variable, then regression analysis is run to determine the best fit and attributes for the stochastic distribution function used to model the lead-times.

Forecast the Combined Demand

Each product has at least one and ideally several different forecasts computed or determined. One of these forecasts is likely to be a manually adjusted forecast. The manually adjusted forecast gives the inventory management team and marketing team the opportunity to put in non-historical factors such as new product introductions, promotional programs, etc. If more than one forecast is used, then a weighted average is used to arrive at the demand forecast to be used in the inventory management processes. Part of the weighting factors includes an assessment as to how accurate the particular forecast methodology has been in the past. The weighting factors are all parameters that are set by the inventory management team. The demand forecast for each product at each distribution point for each period in the forecast horizon is made. For each distribution point, demand is first forecast for customer order fulfillment. Then demand is forecast for tier replenishment based on the forecast data from the lower tier distribution points that replenish from the subject distribution point. The customer fulfillment demand plus the tier replenishment demand together constitute the forecasted demand for the product at the target distribution point.

First, demand to fulfill customer speed and convenience orders is forecast. The mathematical algorithm(s) or heuristic(s) used depends on the settings for the Moving category, the Demand Rate, and the number of units per order. This will determine the forecasted number of orders for speed and convenience. Next, the number of items per order is forecast for the period. If the product is coded as historically experiencing a constant number of items per order, then this constant is used. If the product is coded for lumpy number of items per order, then the fitted stochastic distribution function is used to determine the probability of "k" number of items per order during the period. The value of "k" is determined so that the probability meets a threshold designated by a system parameter. This parameter is closely associated with the desired fill rate. So, for example, the threshold might be set by the inventory management team to be 95%. This establishes the value of "k" such that the probability that the order will contain "k" or fewer items per order is at least 95%. The value of the forecasted items per order is multiplied by the forecasted number of orders to arrive at the forecasted number of speed and convenience demand for the period. The resulting forecast is likely to be a real number (integer plus fraction) rather than a whole integer because the demand is based on stochastic computations.

Next, the forecasted demand for fulfilling ADNs is determined. For this, the system utilizes the Equipment Knowledge Base to analyze the customers who have equipment that might need the product in a maintenance task. These customers are within the service scope of the distribution point such that this

distribution point would fulfill directly to the customer. In Grainger's network, this would be a Branch, a Master Branch or a Distribution Center. Then based on entries in the Equipment Knowledge Base for the condition and age of the equipment and parts together with the reliability stochastic distribution function data, a forecast is made of the probability of needing the product at that customer during the forecast period. This is likely to result in a fractional probability for the individual period. The probabilities of need are added for all the customers in the distribution point's scope. This sum is then rounded to the nearest integer using a threshold parameter as the basis for rounding. This is the threshold that the inventory management team believes is critical in order to comply with the level of service agreements with the customer. For example, this threshold might be specified as 80%. This means that if the probability of need is 80% or more at a distribution point, then we will consider this as a demand for 1 unit for forecasting purposes. The inventory management team considers the 80% likelihood as mandating the physical presence of the product for possible fulfillment. So, in this example, the rounding would occur as follows: 0 to 0.79 would round to zero; 0.8 to 1.79 would round to 1; 1.8 to 2.79 would round to 2; etc. Recall that the distributor will not know if the product is actually needed even when the ADN is submitted for the maintenance work by the customer. The ADN rounding is done separately because the probability of need is different in concept from the probability of demand. The probability of need is based on the age and condition of equipment. It may be 10% in one period, 25% in the next period, 65% in the next period, 90% in the next period, etc. This does not mean that there will be demand for the product in each of the periods. Instead, it means that the likelihood of demand is increasing in each period. The actual demand and need will not be determined until the maintenance team inspects the equipment as part of the maintenance work. It is important that this threshold be implemented as a system parameter since it is likely that the inventory management team will have to modify its value from time to time as experience with the customers is gained. While the rounded amount is used as the forecast to the target distribution point, the unrounded sum is sent up to the next tier for accumulation summing over all the lower tier points. Separate rounding will occur at each tier in the logistics network.

The rounded ADN forecast is then added to the speed and convenience demand computed earlier for the forecast period. Because the speed and convenience demand may still be a real number (whole integer plus a fraction), this sum will also be a real number.

Next, the world factor (if any) is applied to the combined customer fulfillment demand forecast. When world demand if involved, the Markov chain technique is used. With this technique, the world is defined as a group of states and a state-transition diagram is created. Then the probability of moving from one state to another is specified. The inventory management team must determine the current state and state transitions that may occur in each of the forecasting periods. This consummates in the probability that a particular world event state will occur in the forecast period. This event has a correlated impact on the demand (often specified as a percentage increase to be applied). The amount of impact is applied based on the expected probability of the world state occurring.

Now, the tier replenishment forecasted demand is considered. This requires that the demand for all the lower tier distribution points that replenish from the subject distribution point be already computed. The replenishment demand for those lower tiers is added together to create the total forecasted demand for the subject distribution point for the forecast period.

Finally, the historical data will indicate how much of this subject distribution point's inventory is replenished from higher level tier distribution points as oppose to be replenished directly by the supplier. The percent that is replenished through higher level tier distribution points is multiplied by the total forecasted demand to arrive at the forecasted replenishment demand that will be passed up to the higher level distribution point for its forecast computations.

It should be noted that for slow moving product, forecasted demand may only be for the given forecast period fractional (that is either less than 1 or a whole number plus some additional fractional amount such as 2.56). This is especially likely to occur at the lowest tier levels. If this is the situation, then the fractional demand is accumulated over chronological periods until it reaches a quantity greater than or equal to one. This whole unit of forecasted demand is then assigned to the forecast period where the

accumulated fraction first exceeded 50%. Any rounded off residual fractional demand is assigned to the last chronological forecast period for the next chronological accumulation. For example:

Period Forecast	0.15	0.35	0.20	0.45	0.15	0.15	0.20	0.10	0.15	0.25	0.20
Accumulated	0.15	0.50	0.70	1.15	1.30	1.45	1.65	1.75	1.90	2.15	2.35
Forecasted Units	0	1	0	0	0	0	1	0	0	0	0

Another example is:

Period Forecast	1 15	0.35	3.20	0.45	0.15	1.15	0.20	1.10	0.15	0 25	0.20
Accumulated	0 15	1 50	4 7	5 15	5 30	6.45	6.65	7.75	7.90	8.15	8.35
Forecasted Units	1	1	3	0	0	1	1	1	0	0	0

Note that the forecasted number of units of demand for each period is the accumulated amount rounded to the nearest whole integer less the accumulated forecasted units from prior periods. A was done with the ADN fractional amounts, the rounded, whole integers are used for the target distribution point, but the unrounded amounts are forwarded upwards to the next higher tier distribution point when considering replenishment requirements. This enables the system to accumulate fractional amounts and do the rounding at the next tier which may result is a more precise network-wide computation of true demand.

At this point in the process, there is a consolidated forecast of demand for the product at each point in the logistics network.

Determine Critical Stocking Ratios

The critical stocking ratio determines what quantity of product the company can afford to stock in total. It is determined by a mathematical formula that accounts for the product's cost, gross margin, carrying costs, net working asset charges, expected demand, etc. It establishes that maximum amount of product that the company will stock at any one moment in time. The critical stocking ration computations are usually done for a planning horizon of a year since most of the factors going into the function are annualized factors.

Determine Allocations and Base Stocking Levels

A preliminary base stocking level is computed for each distribution point by using the mathematical algorithms and heuristics prescribed by the classification of the product along the Moving-Demand Rate – Number Per Order-World Factor-Lead-time dimensions. These mathematical algorithms and heuristics will also take into account the lead-time data discussed earlier.

Once the preliminary base stocking levels are computed, they are accumulated over the entire logistics network and compared to the maximum stocking level allowed by the critical stocking ratio. If the total is less than the maximum allowed by the critical stocking ratio, then the base stocking levels are finalized and no further adjustments are needed.

If, however, the total exceeds the maximum allowed by the critical stocking ratio, then allocation must take place. If the critical stocking ratio amount will permit at least one unit at each distribution point that has a base stocking level set, then each such distribution point will have at least one unit in its base stocking level. The residual from the critical stocking ratio will then be allocated to those distribution points with the largest preliminary base stocking levels. Once the residual is used up, the remaining distribution points with preliminary base stocking levels will have those levels set to 1.

If the critical stocking ratio does not permit all the distribution points with a preliminary base stocking level to have at least 1 unit stocked, then allocation will be done by giving 1 unit of stock to the distribution points in the order of the distribution points with the largest preliminary stocking levels. The limited

stocking amount is allocated in units of 1 to the distribution points in the sequence of decreasing amounts of preliminary stocking levels.

For example, if the critical stocking ratio allows for 5 units total, then the allocation will look like this:

Inventory Point	A	B	C	D
Preliminary Base Stocking Level	2	2	1	1
Allocated Base Stocking Level	2	1	1	1

However, if the critical stocking ration only allows 3 units, then the allocation will look like this:

Inventory Point	A	B	C	D
Preliminary Base Stocking Level	2	2	1	1
Allocated Base Stocking Level	1	1	1	0

A system parameter is used to permit the inventory management team to have the allocation process give preference to allocating stock to the middle and higher tier distribution points rather than to the lower tier points. This concentrates the limited stocking ability at points that can be quickly dispatched to the lower tiers of the network when the lower tiers experience real demand.

If there is a large difference between the total preliminary base stocking levels and the maximum allowed by the critical stocking ratio, then this situation is brought to the attention of the inventory management team for manual intervention and resolution. There is a system parameter which specifies this difference threshold.

Finally, the inventory management team has the ability to manually override these automatic base level computations if they choose to impose specific stocking models on certain distribution points. For example, certain lowest level tier points (branches in Grainger's situation) may have a list of "never out" items that is maintained for marketing purposes regardless of the financial considerations embraced in the inventory management processes. Or there may be a minimal "national stocking model" that is applied to all branches to have these branches stock a minimal set of common product.

Determine Replenishment Methods

The last step is to determine the replenishment method to be used for the product at each distribution point. The choices are a "use-one-replenish-one" method or a "reorder-point" method where no replenishment is initiated until the unreserved inventory on-hand and in-transit quantity falls below the reorder point. The same mathematical algorithms and heuristics used to determine the preliminary stocking levels will also be used to determine the replenishment methods. In general, most of these algorithms and heuristics determine the choice based on the value of the mean lead-time demand and the relationship between fixed and variable replenishment costs. The mean lead-time demand is the demand anticipated over the lead-time to replenish stock. Lead-time is measured as the time between the point where the replenishment need is computed to the point where the stock has been received and put-away (or cross-docked) and is ready to be used in subsequent fulfillment.

This step concludes the computations and set-up for the inventory management. As mentioned earlier, the entire process is performed periodically so that the forecast horizon keeps moving forward with each forecast iteration.

Run the Initial Replenishment Process

The first replenishment process that is performed after the base stocking levels and re-order points have been set will place the replenishment orders to bring the logistics network into some initial state of readiness. Because of the domino effect in a logistics network topology of many tiers, it may require several replenishment iterations before the inventory state is brought into conformance with the desired state as specified in the inventory planning processes, algorithms and heuristics.

One situation that may develop is that forecasted demand from the current iteration may be substantially less than that forecasted in prior iterations. This may result in on-hand inventory becoming excess. When there is excess inventory, the system will execute a "reverse fulfillment planning process" to determine the disposition of the excess inventory. This process will take into account any future expected increases in demand, costs of handling excess inventory, and candidate consolidation points for the excess. It is quite possible that the process may decide that some excess inventory should simply be left in place. This may be the case if there is still forecasted demand, but the new demand is less than the previous demand. This may also be the case if the costs of moving the excess exceed the value of the excess.

Run Subsequent Replenishment Processes

Subsequent replenishment iterations will take place as the actual demand which was forecasted occurs. By this time, the entire logistics network should be operating in the semblance of "steady state".

The system must capture the accuracy of the forecasts. Accuracy is measured in a couple of ways. One is the amount of overstock that occurs from over forecasting. It will be most apparent in metrics such as total inventory value, inventory turns, amount of excess, etc. Another way is at the other end of the spectrum and is caused by under forecasting. This involves inventory outages and is reflected in order fulfillment rates. The inventory management team uses the historical accuracy of the forecasts to determine the degree to which the mathematical algorithms and heuristics used to create the forecast will be employed in the future or replaced with other alternatives. The accuracy also affects the assignment of weighting factors the inventory management team uses to determine the weighted average forecasted demand.

Correlated Demand

Some products have demand that is correlated to the demand for other products. To determine the impact of correlated demand, the historical orders need to be analyzed to determine what products are bought on the same order as other products. This is exactly the kind of associative buying analysis that retailers such as grocery stores routinely perform using data from the frequent shopper card program.

Because of the large size of the number of products involved, it is necessary to do this analysis many times with each iteration analyzing a manageably sized group of products. There will be a system parameter which specifies the frequency with which this analysis takes place. This frequency then determines the size of the product group (size will be the total size divided by the number of cycles per year – assuming that all products should be analyzed each year).

Customer orders for the last year will be analyzed. For each product in the group being analyzed, a list of correlated product pairs will be created based on what other products were bought on the same order as the target product was bought. The number of occurrences of this pair of purchases will be tallied over the year's worth of order history. A system parameter will specify a threshold of occurrences that must be met in order to consider the product pair to be correlated. This threshold can either be an absolute amount (such as 10 occurrences) or a percentage of the total number of orders involving the target product (such as 60% of all orders of the target product involved also buying the correlated product). Any product pairs failing the threshold test are discarded.

The system then determines the percentage of total demand for the target product that is correlated demand. Product managers will have to provide input as to which of the two products in the pair is the primary (the product purchased independently) and secondary (the one purchased because the other was purchased). This information is then recorded in the product database. When customer speed and convenience demand is computed for the product during the forecasting cycle, it is computed in two parts – the part of demand that comes from primary demand and the part of demand that comes from secondary demand. In order to compute the secondary demand, the demand for the other product in the pair must first be computed. Then the correlation factor from the product database is applied to arrive at the correlated or secondary demand for the target product. The primary and secondary demands are added together to form the total forecasted speed and convenience demand.

Also, when the system determines and analyzes the demand profile for the target product, only primary demand is considered. The demand rate for the secondary demand will follow the demand profile of the other correlated product (which is the primary product of the pair). If the target product is the primary product of the pair, then its entire demand is considered primary and included in the demand profile analysis.

Promotions

Promotions need to be handled manually in the inventory management processes. The marketing team will design and schedule promotions and the inventory management team will work with the marketing team to determine the expected impact of the promotion campaign on demand. This impact must be manually entered into the system as a specific forecast that is considered in the weighted average forecasted demand.

It is also necessary for the system to track the actual results of the promotion so that promotion-generated demand can be excluded from the historical database used in forecasting.

Forecasting Mathematics and Heuristics

Because of the emphasis on uncertainty in the MRO environment, many of the forecasting mathematical algorithms and heuristics will involve stochastic principles and theory. This is one area that differentiates inventory management in the MRO environment from inventory management in other industries.

The factors that generally go into the mathematics and heuristics include cost of product, replenishment costs (both fixed and variable), costs of stockouts which include both lost margins and the potential for further lost business), holding and carrying costs, as well as the stochastic distribution function data.

Often an M/M/M, M/M/G, G/M/M, etc. queuing models can be used to verify or confirm forecasted amounts. Here the queue arrival rate is the order or demand rate. The stochastic distribution function used in the demand rate is the arrival distribution use in the queuing model. The service time is the replenishment lead-time. For variable lead-times, the queuing model uses the same distribution function as is used to model a variable lead-time. The probability of having no customers (orders or demand) in the queue is computed for the queuing model using the coefficient of variation for the particular distribution function. If the probability of no orders in the queue is less than 1, then the inventory point will experience stockouts with a stockout probability of 1 minus the probability of no orders in the queue. Another check is done to determine if the traffic intensity is less than the number of servers (size of the base stock level). If it is not, then stockouts will occur.

Sharing Forecast Data Among Suppliers

One purpose of the integrated supply chain in the MRO environment is to give the suppliers advanced information regarding forecasts of future demand. After each forecast iteration, the suppliers can be given consolidated demand and replenishment forecasts for their products over the forecasting horizon. The consolidated data can be further broken down by anticipated delivery point. Some suppliers may desire the data for the immediate next quarter to be broken down in a finer level of detail such as forecasts by week.

Advance Demand Notice Structure

The occurrance codes are: ? = zero or one (optional), + = one or more, * = zero or more

```
AdvDmdNotice
    AdvDmdNoticeHeader
        Action
        AdvDmdNoticeID
            RequestingEntity
            ADNSerialNumber
        DateTimeStamp
            Date
            Time
        Reason
            ChgReason
            Reason(body)
        StatusReporting
            StatusRq
            StatusRecipient+
                PreferredMedia
                Role
                RecipientName
                RecipientAddress+
                    Media
                    RecipientAddress(body)

    RelationshipParams?
        CustomerRelationship?
            BusUnit
            CustAcctID
            LineOfBusiness?
            PricingMethod?
        CustomerMotivation?
            Init
            PromoID*
        WarrantyOverrides?
            WarrantyTerms*
                Action
                WarrantyTerms(body)
            ReturnPolicy*
                Action
                ReturnPolicy(body)
            MaintPlan?
                Action
                MaintPlan(body)
        SupplierDiversity?
            SupplierDiversity(body)

    ProductParams*
        DemandProbability
            DemandProbability(body)
        DemandQuantity
            ProductQuantity
            ProductQtyUOM
        ProductID
            GraingerSKU?
            CustomerSKU?
            MfgrPartID?
            ProdClassSpec*
                ClassAuth
                ProdClassNode
                ProdClassDesc
                AttrSpec*
                    CanBeChanged
                    ChangePriority
                    ClassAttr
                    AttrValue+
                    MinValue?
                    MaxValue?
```

```
                ProdClassSpec*
                ProdSubRule?
                ApprovalPolicyRule?
            ProdName*
            CompetitorSKU?
            ApplicationUsage?
                MaintenanceTask?
                    MaintTaskNode
                    MaintTaskDesc
                    EquipmentID?,
                    MaintenanceTask*
                EngineeringTask?
                    EngrTaskNode
                    EngrTaskDesc
                    EngrTaskParam*
                    EngineeringTask*
                BOMBuild?
                    SuppLst
                    PartsLst
                    WorkKit
                    GraingerSKU?
                    EquipmentID?
            WorkKit
            Exactness
            UnitShip
            UnitIssue
            BarCode
            MSDS
            OprInst
            AssbyInst
            ApplNotes
            PMPlans
            Training
        AgeReplacing?
            AgeUOM
            AgeMetric

    ServiceParams*
        DemandProbability
        ServiceID
            ServiceCatg
            ServiceName?

    LogisticsParams?
        Situatn
        StatusRq
        LogisticsMode
            PickUp
                BranchNumber
            ShipToCustomer
                DeliveryRequirements
                    WhereNeeded
                        StreetAddress+
                        CityName
                        StateProvince
                        ZIPPostalCode
                        Country
                        LocDUNSCode?
                    InternalDeliveryPt?
                    DeliveryProcedure?
                    DeliveryContact*
                        PersonName
                        PhoneNumber+
                        EmailAddress?
                ShipOverride?
                    ShipMethod?
                        PkgReq
                        ShipMethod(body)
                    Carrier?
                        CarrierService
                        BillToCarrierAcct
```

```
                                        Carrier(body)
                                ShipCostLimit?
                                        CurrencyAmt
                                        CurrencyCountry?
                                FOBPoint?
                LOSHold
                        GraingerFacilityNumber
                        ReleaseNotNeededDate
        WhenNeeded?
                EarlistNeeded?
                LatestNeeded
        ShipLabelReq?
        Equipment*
                Action
                EquipmentID
                        Manufacturer
                        MfgrPartNo
                        RevisionNumber?
                        SerialNumber
                EquipmentInstance?
                MaintenanceProcedure?

TransactionParams?
        TransType
        FormPay
        CustTransRef?
```

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!-- rev: 05-03-2000 rhs to add age, location DUNS code, and equipment revision/serial no. -->

<!DOCTYPE AdvDmdNotice [
<!ELEMENT AdvDmdNotice  (AdvDmdNoticeHeader,RelationshipParams?,ProductParams*,
        ServiceParams*,LogisticsParams?,TransactionParams?) >

<!--The Advance Demand Notice Header provides additional header data that is unique
to this transaction type.-->
<!ELEMENT AdvDmdNoticeHeader
(AdvDmdNoticeID,DateTimeStamp,Reason,StatusReporting) >

<!ATTLIST AdvDmdNoticeHeader
        Action (New | Delete | Revision)  #REQUIRED  >

<!--This section describes the type of relationship that the customer has with Grainger.-->
<!ELEMENT RelationshipParams  (CustomerRelationship?,CustomerMotivation?,
        WarrantyOverrides?,SupplierDiversity?) >

<!--There will be one set of product parameters for each product within the domain of the
demand object. A product is removed from the demand object by setting its quantity
to zero or its probability to zero.-->
<!ELEMENT ProductParams  (DemandProbability,DemandQuantity,ProductID,AgeReplacing?)
>

<!--There will be one set of service parameters for each service to be included with this
demand. A service is removed from the demand object by setting its probability to
zero.-->
<!ELEMENT ServiceParams  (DemandProbability,ServiceID) >

<!--Situatn is the situation attribute. Allowed values are:

ConsumeRepl = replenishing a consumable
ConvnASAP /= ordinary speed and convenience
ConvnEmerg = speed and convenience responding to an emergency
DisasterRecv = responding to disaster recovery situations
Disposal = a reverse demand where the customer needs to dispose of an old or used
item.
EquipInstall = demand is needed as part of an equipment installation effort
JIC = demand is to create Just-In-Case safety inventory for the customer
MaintCheck = product is needed to perform a maintenance check on some equipment
PrevMaint = product is needed to perform a preventive maintenance task
SafetyInsp = product is needed to perform a safety inspection
Surplus = a reverse demand where the customer needs toreturn surplus items.-->
<!ELEMENT LogisticsParams  (LogisticsMode,WhenNeeded?,ShipLabelReq?,Equipment*) >

<!ATTLIST LogisticsParams
        Situatn (ConsumeRepl | ConvnASAP | ConvEmerg | DisasterRecv |
        EquipInstall | JIC | MaintCheck | PrevMaint | SafetyInsp | Surplus)
         #REQUIRED
        StatusRq (Bi-Weekly | Daily | Events | Hourly | Monthly | Weekly)
        "Events" >

<!--This section describes how the customer plans to turn this demand into a legal order.
```

```
If the type is a demand advice, then some other form of payment will be used to create
the official order. If the type shown is an order, then this demand object becomes
the legal order.

TransType designates the legal status of this demand object. Demand advices need
some other form of legal ordering such as a PO.-->
<!ELEMENT TransactionParams  (CustTransRef?) >

<!ATTLIST TransactionParams
        TransType (DmdAdvice | Order)  #REQUIRED
        FormPay (BPORelease | Cash | CreditCard | EDI | Invoice | PO |
        PurchCard)  #REQUIRED >

<!--The Advance Demand Notice ID provides a unique system-wide serial number for
the demand object. All event activity for this object uses this ID.-->
<!ELEMENT AdvDmdNoticeID  (RequestingEntity,ADNSerialNumber) >

<!ELEMENT DateTimeStamp  (Date,Time) >

<!--The reason for the change is described in this section.

ChgReason codes are: Insp = Demand data has changed because of inspection
of the equipment; Schd = Demand data has changed because of a change in the
customer's operating or maintenance schedules; Other = Demand data is changing
because of some other reason.-->
<!ELEMENT Reason  (#PCDATA) >

<!ATTLIST Reason
        ChgReason (Insp | Schd | Other)  #REQUIRED
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!--Status Reporting specifies when fulfillment status update information should be passed
to the originator. This element applies to the entire ADN.-->
<!ELEMENT StatusReporting  (StatusRecipient+) >

<!ATTLIST StatusReporting
        StatusRq (Bi-Weekly | Daily | Events | Hourly | Monthly | Weekly)
        "Events" >

<!ELEMENT CustomerRelationship  (BusUnit,CustAcctID,LineOfBusiness?,PricingMethod?) >

<!--Init values are: Contract = the demand resulted from a contractual relationship;
Promo = customer is responding to promotional activity;
Routine = demand comes from routine maintenance work where Grainger is the usual
supplier;
SalesCall = demand results from a sales call on the customer;
WalkIn = the customer was a walk-in at the branch
WWW = the customer came through an Internet channel (Grainger.Com, OrderZone.com)
ProcureSys = the demand came from the customer via a procurement system such
as Ariba or CommerceOne;
EComm = the demand came from the customer via some other electronic commerce
channel.
-->
```

<!ELEMENT CustomerMotivation (PromoID*) >

<!ATTLIST CustomerMotivation
Init (Contract | Promo | Routine | SalesCall | WalkIn | WWW |
ProcureSys | EComm) #REQUIRED >

<!--Normally these aspects of relationships come from either company policy defaults or from specifics contract with the customer and kept in the customer data base. This section is used only if the defaults are to be overridden. They are included in the demand object since they may be a factor in determining a price to the customer. To delete a section, the appropriate attributes are set.-->
<!ELEMENT WarrantyOverrides (WarrantyTerms*,ReturnPolicy*,MaintPlan?) >

<!--Allows the customer to specify any supplier diversity requirements for consideration in fulfilling this demand.-->
<!ELEMENT SupplierDiversity (#PCDATA) >

<!ATTLIST SupplierDiversity
LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
ISO8061Date | Time-msecZ) "Char"
Action (New | Delete | Revision) #REQUIRED >

<!--Setting the demand probability to 100% essentially treats the demand like a firm order.-->
<!ELEMENT DemandProbability (#PCDATA) >

<!ATTLIST DemandProbability
LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
ISO8061Date | Time-msecZ) "Float" >

<!ELEMENT DemandQuantity (ProductQuantity,ProductQtyUOM) >

<!--The product is identified by one of the methods. Note that subsequent analysis may indicate that product classification and application usage may be done prior to demand creation so that demand creation actually knows the products involved.

Attributes:
WorkKit = Yes/No attribute to include work kit items as well as the basic product. Determine the tools and ancillary parts needed when doing work on the primary part.
Exactness = Indicates the degree of exactness that is required in identifying the product. Allowed values are: ExactRpl when the exact replacement is needed; PeerRpl when a peer equivalent (such as a similar part made by a different manufacturer) is permitted if an exact replacement is not found; UpGrdRpl when a substitute product of higher capability is permitted if either the exact match or peer replacement cannot be found.
UnitShip = the packaging aggregation the customer requires for shipping.
UnitIssue = the packaging aggregation the customer requires for issueing to end-users.
BarCode = any bar coding requirements the customer has
MSDS = a Yes/No attribute indicating whether an MSDS sheet should be included with the product.
OprInst = a Yes/No attribute indicating if operating instructins should be included with the product.
AssbyInst = a Yes/No attribute indicating if assembly instructions should be included with the product.

```
ApplNote = a Yes/No attribute indicating if application notes should be included
with the product.
PMPlans = a Yes/No attribute indicating if PM plans should be included with the
product.
Training = a Yes/No attribute indicating if training materials should be included with
the product.-->
<!ELEMENT ProductID
(GraingerSKU?,CustomerSKU?,MfgrPartID?,ProdClassSpec*,ProdName*,
        CompetitorSKU?,ApplicationUsage?) >

<!ATTLIST ProductID
        WorkKit (Yes | No) "No"
        Exactness (ExactRpl | PeerRpl | UpGrdRpl)  #REQUIRED
        UnitShip  CDATA   #REQUIRED
        UnitIssue  CDATA   #REQUIRED
        BarCode  CDATA   #IMPLIED
        MSDS (Yes | No)  "No"
        OprInst (Yes | No)  "No"
        AssbyInst (Yes | No)  "No"
        ApplNotes (Yes | No)  "No"
        PMPlans (Yes | No)  "No"
        Training (Yes | No)  "No" >

<!ELEMENT AgeReplacing   EMPTY  >

<!--This gives the age of the current part that is being replaced by the part in the ADN.
This age is used to determine the probability of needing to actually replace the part
when the maintenance action is completed.

The AgeUOM uses one of the Engineering Unit Table entries from the MIMOSA standard
as specified in the Equipment Knowledge Base.-->
<!ATTLIST AgeReplacing
        AgeUOM  CDATA   #REQUIRED
        AgeMetric  CDATA   #REQUIRED  >

<!--Some scheme for identifying services still needs to be developed.-->
<!ELEMENT ServiceID  (ServiceCatg,ServiceName?) >

<!--Indicates is this is a customer pick-up, a customer ship order, or held for a level of
service agreement with the customer.-->
<!ELEMENT LogisticsMode  (PickUp | ShipToCustomer | LOSHold) >

<!ELEMENT WhenNeeded  (EarlistNeeded?,LatestNeeded) >

<!ELEMENT ShipLabelReq  (#PCDATA) >

<!ATTLIST ShipLabelReq
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!--This field is used to associate the products and services requested in the demand to
maintenance on a particular type and piece of equipment. This data will be used to
update the maintenance knowledge base.-->
<!ELEMENT Equipment  (EquipmentID,EquipmentInstance?,MaintenanceProcedure?) >
```

```
<!ATTLIST Equipment
        Action (New | Delete | Revision)  #REQUIRED  >

<!--This field is used to specify PO number, credit card number, etc.-->
<!ELEMENT CustTransRef  (#PCDATA) >

<!ATTLIST CustTransRef
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!--The Requesting Entity uniquely identifies who is originating this demand. Requesting
Entity identifiers are assigned by the System Administrator.-->
<!ELEMENT RequestingEntity  (#PCDATA) >

<!ATTLIST RequestingEntity
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!--The ADN Serial Number is unique to the requesting eneity. The requesting entity
assigns this serial number.-->
<!ELEMENT ADNSerialNumber  (#PCDATA) >

<!ATTLIST ADNSerialNumber
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "UnsignedLong" >

<!ELEMENT Date  (#PCDATA) >

<!ATTLIST Date
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "ISO8061Date" >

<!ELEMENT Time  (#PCDATA) >

<!ATTLIST Time
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Time-msecZ" >

<!ELEMENT StatusRecipient  (RecipientName,RecipientAddress+) >

<!ATTLIST StatusRecipient
        PreferredMedia (Phone | Email | Pager | Fax | WebPage)  "Email"

        Role (EndUser | MaintMgr | Purchaser | MROMgr | GraingerSalesTeam)
         #REQUIRED  >

<!--This is the W.W. Grainger, Inc. business unit (such as Grainger Industrial Supply, etc)-->
<!ELEMENT BusUnit  (#PCDATA) >

<!ATTLIST BusUnit
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >
```

```
<!--Customer Account ID's can be unique within a business unit.-->
<!ELEMENT CustAcctID (#PCDATA) >

<!ATTLIST CustAcctID
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!--The line of businessis used only if the business unit needs to track along lines of business or
or if the line of business determines other customer relationship parameters to use.-->
<!ELEMENT LineOfBusiness (#PCDATA) >

<!ATTLIST LineOfBusiness
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!--The pricing method is used t override the standard, default customer-based pricing
methor normally used for the customer.-->
<!ELEMENT PricingMethod (#PCDATA) >

<!ATTLIST PricingMethod
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT PromoID (#PCDATA) >

<!ATTLIST PromoID
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT WarrantyTerms (#PCDATA) >

<!ATTLIST WarrantyTerms
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char"
        Action (Change | Remove) "Change" >

<!--The return policy describes the circumstances under which a customer will be allowed
to return product associated with this demand and also specifies any return or
restocking fees to be assessed.-->
<!ELEMENT ReturnPolicy (#PCDATA) >

<!ATTLIST ReturnPolicy
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char"
        Action (Change | Remove) "Change" >

<!ELEMENT MaintPlan (#PCDATA) >

<!ATTLIST MaintPlan
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char"
        Action (Change | Remove) "Change" >

<!ELEMENT ProductQuantity (#PCDATA) >
```

```
<!ATTLIST ProductQuantity
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "UnsignedLong" >

<!ELEMENT ProductQtyUOM  (#PCDATA) >

<!ATTLIST ProductQtyUOM
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT GraingerSKU  (#PCDATA) >

<!ATTLIST GraingerSKU
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT CustomerSKU  (#PCDATA) >

<!ATTLIST CustomerSKU
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT MfgrPartID  (Manufacturer,MfgrPartNo) >

<!--The product classification is a hierarchy, and the complete product specification must
include ass applicable levels of the classification hierarchy.

The ClassAuth indicates whose product classification hierarchy is being used.-->
<!ELEMENT ProdClassSpec  (ProdClassNode,ProdClassDesc,AttrSpec*,ProdClassSpec*,
        ProdSubRule?,ApprovalPolicyRule?) >

<!ATTLIST ProdClassSpec
        ClassAuth (Customer | DunBrad | Grainger)  "Grainger"
        >

<!--The product name must be a market-recognized, brand-like name which uniquely
identifies a product.-->
<!ELEMENT ProdName  (#PCDATA) >

<!ATTLIST ProdName
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT CompetitorSKU  (CompetitorID,CompetitorPartNo) >

<!--This allows the products to be selected from a description of how they will be used.-->
<!ELEMENT ApplicationUsage  (MaintenanceTask?,EngineeringTask?,BOMBuild?) >

<!ELEMENT ServiceCatg  (#PCDATA) >

<!ATTLIST ServiceCatg
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >
```

```
<!ELEMENT ServiceName  (#PCDATA) >

<!ATTLIST ServiceName
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!--The branch number must be a valid, Grainger branch number.-->
<!ELEMENT PickUp   EMPTY  >

<!ATTLIST PickUp
        BranchNumber  CDATA   #REQUIRED  >

<!ELEMENT ShipToCustomer  (DeliveryRequirements,ShipOverride?) >

<!--The value of LOSHold references the Level of Service contract between Grainger
and the customer that applies to this demand situation.
The Level of Service Hold location is where the level of service agreement with the
customer specifies the products will be staged pending the customer's determination
of actual need. The Grainger Facility Number specifies either a branch number or a
distribution center number where the products are to be staged. The Release Not
Needed Date indicates when Grainger is to assume that the products are not
needed and can release them for commitment to another demand or order.-->
<!ELEMENT LOSHold  (#PCDATA) >

<!ATTLIST LOSHold
        GraingerFacilityNumber  CDATA   #REQUIRED
        ReleaseNotNeededDate  CDATA   #REQUIRED  >

<!ELEMENT EarlistNeeded  (DateSpec) >

<!ELEMENT LatestNeeded  (DateSpec) >

<!ELEMENT EquipmentID  (Manufacturer,MfgrPartNo,RevisionNumber?,SerialNumber?) >

<!--When a customer has multiple instances of the same type of equipment, this field is
used to indicate which instance is involved in the maintenance.-->
<!ELEMENT EquipmentInstance  (#PCDATA) >

<!ATTLIST EquipmentInstance
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!--Identifies the maintenance procedure being performed on the equipment which is
generating the demand.-->
<!ELEMENT MaintenanceProcedure  (#PCDATA) >

<!ATTLIST MaintenanceProcedure
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!--The actual name of the status recipient.-->
<!ELEMENT RecipientName  (#PCDATA) >
```

```
<!ATTLIST RecipientName
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!--The recipient's address according to the media type. If the media is specified as a
Web Page, then the recipient will go to the Grainger web site to access the status
rather than having the status pro-actively sent to him/her.-->
<!ELEMENT RecipientAddress (#PCDATA) >

<!ATTLIST RecipientAddress
        Media (Phone | Email | Pager | Fax | WebPage) #REQUIRED
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT Manufacturer (#PCDATA) >

<!ATTLIST Manufacturer
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT MfgrPartNo (#PCDATA) >

<!ATTLIST MfgrPartNo
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT ProdClassNode (#PCDATA) >

<!ATTLIST ProdClassNode
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT ProdClassDesc (#PCDATA) >

<!ATTLIST ProdClassDesc
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!--The CanBeChanged indicates if this attribute can be changed in order to find a
suitable substitute/replacement product if the target product is not available. The
Change Priority indicates the priority sequence in which this attribute is changed relative
to the other attributes that can be changed. A priority of 1 indicates the first attribute
that will be changed, 2 the second, and so forth.-->
<!ELEMENT AttrSpec (ClassAttr,AttrValue+,MinValue?,MaxValue?) >

<!ATTLIST AttrSpec
        CanBeChanged (Yes | No) "No"
        ChangePriority CDATA #REQUIRED >

<!--The product substitution rule contains a reference to the rules that must be followed
when seeking a substitute product if the target product is unavailable.-->
<!ELEMENT ProdSubRule (#PCDATA) >

<!ATTLIST ProdSubRule
```

```
          LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
          ISO8061Date | Time-msecZ) "Char" >

<!--Contains a referent to the rules (in the rules system/server) that must be followed when
seeking a substitute product if the target product is unavailable.-->
<!ELEMENT ApprovalPolicyRule  (#PCDATA) >

<!ATTLIST ApprovalPolicyRule
          LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
          ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT CompetitorID  (#PCDATA) >

<!ATTLIST CompetitorID
          LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
          ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT CompetitorPartNo  (#PCDATA) >

<!ATTLIST CompetitorPartNo
          LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
          ISO8061Date | Time-msecZ) "Char" >

<!--Maintenance tasks are organized into a classification hierarchy. The identification of
a maintenance task requires the complete hierarchical description.-->
<!ELEMENT MaintenanceTask  (MaintTaskNode,MaintTaskDesc,EquipmentID?,
          MaintenanceTask*) >

<!--Engineering tasks are situations where a particular product is needed based on its
ability to meet the requirements for whatever is being designed/engineered. Engineering
tasks are classified hierarchically.-->
<!ELEMENT EngineeringTask  (EngrTaskNode,EngrTaskDesc,EngrTaskParam*,
          EngineeringTask*) >

<!--The BOM Build is used to determine what other parts are needed in conjunction with
a primary part of piece of equipment.

SuppLst is a Yes/No attribute indicating whether the operational supplies are needed
for the primary part.
PartsLst is a Yes/No attribute indicating whether the spare parts for the primary part
are needed.
WorkKit is a Yes/No attribute indicating whether other parts in a work kit are needed.-->
<!ELEMENT BOMBuild  (GraingerSKU?,EquipmentID?) >

<!ATTLIST BOMBuild
          SuppLst (Yes | No)  "No"
          PartsLst (Yes | No)  "No"
          WorkKit (Yes | No)  "No" >

<!ELEMENT DeliveryRequirements  (WhereNeeded,InternalDeliveryPt?,DeliveryProcedure?,
          DeliveryContact*) >

<!ELEMENT ShipOverride  (ShipMethod?,Carrier?,ShipCostLimit?,FOBPoint?) >
```

```
<!ELEMENT DateSpec (Date,Time?) >

<!ELEMENT RevisionNumber (#PCDATA) >

<!ELEMENT SerialNumber (#PCDATA) >

<!ELEMENT ClassAttr (#PCDATA) >

<!ATTLIST ClassAttr
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char"
        ValType (Bool | Enum | LongDouble | UnsignedLong) #REQUIRED >

<!--The attribute value is the preferred, desired value of the attribute.-->
<!ELEMENT AttrValue (#PCDATA) >

<!ATTLIST AttrValue
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!--The minimum attribute value to which the attribute can be changed when seeking an
alternative.-->
<!ELEMENT MinValue (#PCDATA) >

<!ATTLIST MinValue
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!--The maximum attribute value to which the attribute can be changed when seeking an
alternative.-->
<!ELEMENT MaxValue (#PCDATA) >

<!ATTLIST MaxValue
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT MaintTaskNode (#PCDATA) >

<!ATTLIST MaintTaskNode
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT MaintTaskDesc (#PCDATA) >

<!ATTLIST MaintTaskDesc
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT EngrTaskNode (#PCDATA) >

<!ATTLIST EngrTaskNode
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >
```

```
<!ELEMENT EngrTaskDesc  (#PCDATA) >

<!ATTLIST EngrTaskDesc
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!--The engineering parameters needed to determine what products are needed.-->
<!ELEMENT EngrTaskParam  (#PCDATA) >

<!ATTLIST EngrTaskParam
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT WhereNeeded  (StreetAddress+,CityName,StateProvince,ZIPPostalCode,Country,
        LocDUNSCode?) >

<!ELEMENT InternalDeliveryPt  (#PCDATA) >

<!ATTLIST InternalDeliveryPt
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!--This references the customer's delivery procedure to be used.-->
<!ELEMENT DeliveryProcedure  (#PCDATA) >

<!ATTLIST DeliveryProcedure
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT DeliveryContact  (PersonName,PhoneNumber+,EmailAddress?) >

<!--Ship method examples are Ground, Overnight, Air, etc.

The PkgReq specifies any special packaging requirements needed for the method of
shipment such as hazardous materials packaging, special labeling, etc..-->
<!ELEMENT ShipMethod  (#PCDATA) >

<!ATTLIST ShipMethod
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char"
        PkgReq  CDATA   #IMPLIED  >

<!--The Carrier Service is the particular service type to be used (for example, overnight or
2-day ground) and the Bill To Carrier Account number is the account number with the
carrier to which the shipment is billed.-->
<!ELEMENT Carrier  (#PCDATA) >

<!ATTLIST Carrier
        CarrierService  CDATA   #REQUIRED
        BillToCarrierAcct  CDATA   #REQUIRED
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT ShipCostLimit  (CurrencyAmt,CurrencyCountry?) >
```

```
<!ELEMENT FOBPoint  (#PCDATA) >

<!ATTLIST FOBPoint
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT StreetAddress  (#PCDATA) >

<!ATTLIST StreetAddress
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT CityName  (#PCDATA) >

<!ATTLIST CityName
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT StateProvince  (#PCDATA) >

<!ATTLIST StateProvince
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT ZIPPostalCode  (#PCDATA) >

<!ATTLIST ZIPPostalCode
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT Country  (#PCDATA) >

<!ATTLIST Country
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!–The DUNS code for the location is used to associate the equipment's location in
the Equipment Knowledge Base.–>
<!ELEMENT LocDUNSCode  (#PCDATA) >

<!ELEMENT PersonName  (#PCDATA) >

<!ATTLIST PersonName
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT PhoneNumber  (#PCDATA) >

<!ATTLIST PhoneNumber
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ)  "Char" >

<!ELEMENT EmailAddress  (#PCDATA) >
```

```
<!ATTLIST EmailAddress
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT CurrencyAmt  (#PCDATA) >

<!ATTLIST CurrencyAmt
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

<!ELEMENT CurrencyCountry  (#PCDATA) >

<!ATTLIST CurrencyCountry
        LexType (Char | Bool | Enum | LongDouble | UnsignedLong | Float |
        ISO8061Date | Time-msecZ) "Char" >

]>

-------------------------76912060530050--
```

MRO ISCM Intelligent Order Fulfillment Planning for S&C Orders and ADN's

Intelligent Order Fulfillment Planning Technical Architecture Approach

The overall approach to intelligent order fulfillment planning is to construct a list of alternative ways to fulfill the order and then to grade and rank each alternative. The "best" alternative is then selected and a fulfillment plan is for that alternative is attached to the order.

A branch and bound technique is used to determine candidate sourcing points for each line item of the order. This approach reduces the number of alternatives that might otherwise be considered by using every combination and permutation of sourcing points throughout the logistics network. The branch and bound approach is implemented by selecting tiers of sourcing points. The exact criteria for defining each tier depends on the specific logistics network topology that is being used. Therefore a table and parameter driven approach is used to define these tiers and how to assign sourcing points to each tier. The example in this document discusses how the sourcing tiers for the Grainger logistics network as it existed during the winter, 2001 would occur.

The logic followed within each tier group of sourcing points will be specified using a rules-based system approach. This allows the fulfillment planning system to be used for any kind of logistics network topology. As with the tier definition, the example in this document discusses how the alternative creation process rules would result in the environment of Grainger's logistics network topology and business model during the Winter, 2001.

By maximizing the use of parameter tables and rules, the basic intelligent fulfillment planning process can be used in any kind of logistics network topology. This is important because as a company changes its logistics network topology to respond to changes in business activity and sales patterns, the fulfillment planning process needs to be adapted quickly without major programming required.

Plan Initiation

When an Order or ADN is received at this point of the order processing process, the following pieces of data will be used for the initial intelligent order fulfillment planning iteration:

From the Order/ADN Header:

- Type of Order (Walk-In, Will-Call, Ship, ADN Reserve)
- Point of Delivery (for ADN the Level of Service will be the be customer's delivery site designated by the ZIP code or the code of the Grainger facility where the ADN is to be staged per the Level of Service agreement, for Walk-In and Will-Call, the code for the Grainger Branch that normally services this customer's site, and for Ship, the code for the Grainger Distribution Center that supports the customer's normal branch.)
- Code for the Branch that normally services this customer's site
- Customer's end-delivery site ZIP code (used for orders and ADN's that will be shipped to the customer)
- The Delivery Date (for ADN's, the date on which the product(s) is/are to be staged at the LOS point, ship orders and will-call orders will show the date the customer wants the product either shipped or at the will-call counter, and for walk-in orders, will show the current date)
- Whether or not the customer requires a single consolidated shipment or whether the customer will accept multiple split shipments.

For Each Line Item:

- The Grainger SKU of the part
- The quantity ordered

- Any supplier diversity requirements for the line item

Determine Product Equivalents

For each line item product, if any equivalent products are in the catalog, they are designated for the line item. An equivalent product is one that has the same function/feature attributes as the ordered product. For example, the catalog shows that 1H385, a NEMA1 Fractional Horsepower Manual Starter has an equivalent alternative of 1H386, the same unit except with a red pilot light. This determination is based on how the products are defined in the system's product catalog database. If there are any supplier diversity requirements specified by the customer, these will be taken into account when looking for equivalent products.

Determine Primary and Tier 1 Secondary Sourcing Points

The primary sourcing point is the Grainger branch normally servicing the customer's delivery site for will-call orders, the Grainger branch which is placing the walk-in order, the Grainger distribution center assigned to the customer site's normal branch from which Ship-To orders will be fulfilled, the Grainger facility where an ADN product is to be staged, or the Grainger distribution center assigned to the customer's site from which an ADN product is shipped for staging at the customer's site. Certain items are only stocked at the National Distribution Center (designated in the product catalog database). In this case, the primary sourcing point will be the NDC. Other items are only handled as direct drop ships from the suppliers (designated in the product catalog database). If direct drop ships are involved, then fulfillment plans for those line items are set up for the drop ship process. This type of logic is parametrically described using a rules-based approach. The assignment of branches to customers is done in the customer master database, and the assignment of distribution center to branches is done in the distribution logistics network database.

For walk-in and will-call orders, the tier 1 secondary sourcing points will be other Grainger branches within a 25-mile radius of the primary sourcing branch. Where the primary sourcing point is a branch for ADN staging orders, the tier 1 secondary source will be the distribution center servicing that branch. Where the primary sourcing point is a distribution center, the tier 1 sourcing points will be other distribution centers which can ship either to the customer's site (for Ship-To orders and ADN's that are to be staged at the customer's site) or can ship to the primary sourcing distribution center (for ADN's that are to be staged at the primary sourcing point distribution center). These tier 1 distribution centers must be able to move the product to the customer or the primary sourcing distribution center before the designated order Delivery Date using the normal shipping mode (usually LTL) between the two points or using Package Shipment to the customer's site. Parameter tables are used to define criteria such as the miles-radius for finding proximity branches and shipping time tables for determining candidate tier 1 distribution centers. Rules-based parameters are used to find the tier 1 candidates that meet the parameter table criteria.

Primary and tier 1 secondary sourcing points will be assigned to each line item.

When considering tier 1 secondary sourcing points, sourcing points are eliminated if they cannot handle direct shipment to customer site's for Ship-To's and ADN's staged at the customer site. They are eliminated if they are not capable of shipping Hazmat. Some locations also have restrictions on categories of products (such as paper products only being shipped from a limited number of distribution centers). In the case of will-call's, the tier 1 secondary sourcing point must be capable of having the customer come for the will-call. Walk-in orders are converted to will-call orders if the customer is referred to a tier 1 secondary sourcing point for fulfillment.

MRO ISCM Intelligent Order Fulfillment Planning for S&C Orders and ADN's

Determine Inventory Availability at the Primary and Tier 1 Sourcing Points

For each line item, determine the on-hand inventory at the primary and tier 1 sourcing points for both the product ordered and for any alternative equivalent products. The on-hand quantity is subdivided according to that which is reserved for other orders and that which is unreserved. If there is in-transit inventory that is expected to arrive and be put-away or cross-docked in time to permit an on-time pick/pack and deliver of the order, then this in-transit inventory can also be included in the on-hand inventory amounts. This type of logic is also implemented using rules-based approaches. Specific parameters such as put-away and cross-dock times are kept in parameter tables so they can be easily changed as the logistics network environment changes.

First Iteration Plan

If the entire order or ADN can be filled by using unreserved on-hand inventory at the primary sourcing point, that create a fulfillment plan to do so. This will be the only fulfillment plan alternative.

If not, then generate the first iteration of fulfillment alternatives to consider.

First, determine if currently reserved on-hand inventory at the primary sourcing point can feasibly be reallocated to the subject order. This can be done if the replenish time for that product will still be done in time to allow the receiving and put-away plus the pick/pack and deliver to be accomplished on time for the currently reserved order. The costs of replenishment, however, will be assigned to the current subject order, not the order currently holding the inventory reservation since it is the subject order that is causing the replenishment requirement. The alternatives created from this consideration are added to the list to be evaluated.

Next, try to fulfill the entire order at a tier 1 secondary sourcing point. If there are still no fulfillment plan alternatives, try the reallocation of reserved inventory as was done for the primary sourcing point. The same criteria and cost allocation is used. (A candidate fulfillment plan alternative requires that there be a fulfillment sourcing point for each line item to completely fill the quantity of that line item. The line item can be split between two sourcing points if needed.)

Next, create the permutation/combination of fulfilling either entire or partial line items at the primary and tier 1 secondary sourcing points.

If there are now one or more candidate fulfillment alternatives cost and evaluate each alternative. Select the best alternative and create the fulfillment plan for the order/ADN accordingly.

However, if there are still no candidate alternatives, then go on to the Second Iteration Plan.

Costing and Evaluating Fulfillment Alternatives

This costing process will be used whenever there are candidate fulfillment alternatives that need to be costed and evaluated. The steps involved are routing each line item from sourcing point to delivery point, determining the activity costs for each line item's plan, considering and evaluating other factors, computing a total weighted grade for each alternative and then selecting the "best" alternative.

After the "best alternative" is selected, it is established as "the" fulfillment plan for the order/ADN and inventory reservations are made accordingly. If reservation re-allocations are involved, appropriate replenishment orders are generated and reserved for the pre-empted order/ADN. New fulfillment plans for those pre-empted orders/ADN's are created to reflect the new replenishment activity.

Routing the Line Items

How each line item is routed depends on whether the customer requires a single consolidated shipment. If consolidation is required, then additional candidate alternatives must be generated to consider candidate consolidation points. Three candidate consolidation points will be considered:

- Consolidate at the point where the most line items are being sourced.
- Consolidate at the point that minimizes the shipment of all items not sourced at that point to the consolidation location.
- Consolidate at the point that minimizes the cost of final shipment of the consolidated order to the customer or the ADN staging point.

Each non-routed fulfillment alternative will be expanded into three routed alternatives, one for each of the possible consolidation points.

If the customer does not require a consolidated shipment, it may still be advantageous to consolidate. This situation might occur if there are multiple line items with small-sized and/or low-valued items. This is because the Package Shipment shipping costs to the customer of the multiple items may exceed the combined shipping costs of moving the product using LTL to a consolidation point and then using Package Shipment for the final shipment to the customer. Since this trade-off is very situation-dependent, the only way to objectively determine it is to create the consolidation alternatives and then compare them to the non-consolidation alternatives. A parameter table is used to define the size and value thresholds. If this situation occurs, then consolidation alternatives will be generated even if the customer does not require them.

Determining Activity Costs for Each Line Item

The activity costs are accumulated as the product movement is traced from the point of sourcing to the point of final delivery or ADN staging. The activity costs include:

- Picking costs every time the product is picked
- Put-away costs every time the product is put away
- Cross-docking costs is this occurs en-route to consolidation
- Carrying costs if the product is inventoried temporarily during the consolidation process
- LTL shipping costs between Grainger logistics points
- Package Shipment shipping costs between Grainger logistics points if faster shipping is required to meet order delivery deadlines than can be achieved with the LTL shipping
- Package Shipment shipping costs to the customer
- Holding costs of the product's value over the delivery fulfillment period

The activity costs are kept in parameter tables by location and by the physical attributes of the product (size, weight, Hazmat, etc.) The parameter values to be used in computing the activity costs are kept in the product catalog database.

Considering and evaluating Other Factors

Parameter tables are used to determine which of the following other factors are to be considered in evaluating the candidate alternatives. The use of these factors can vary depending on the current situation of the business.

Opportunity costs can be considered over the period that it takes to replenish the product. If there are other known orders and ADNs for other customers, then there is an opportunity cost associated with using the

candidate product for the current order and not fulfilling the order for the other customer. The magnitude of this opportunity can be partially based on the relative preferences of the two customers. In the case of ADNs, the difference between the probability of need for each order/ADN needs to be considered. (Regular orders are considered to have a 100% probability of need.) Another form of opportunity cost that can be considered is an anticipated peak in demand for the product at the sourcing location. This can occur either because of seasonal changes or because recent demand has been significantly below expectation and the law of averages expects that the short-term demand will be larger in order to reach the overall average demand over time. In this case, the computed opportunity cost will tend to favor alternatives where product is sourced from a point not anticipating a short-term increase in demand over a point that is anticipating such a demand.

Another factor that can be considered is the degree to which the alternative uses excess inventory. Excess is defined as the amount of unreserved inventory on hand that exceeds the stocking level for the stocking location. Excess is also considered if the product has not had a demand hit at the particular location for the last "x" months (the value of "x" is a parameter which can be set on a location-specific basis or can be set for the overall company). If the use of excess inventory will be a factor in the evaluation of the alternative, then the "cost" value has to be negative in excess situations in order to have the evaluation favor the excess over the non-excess alternative. This can be done by setting the cost equal to the negative of the carrying cost for "x/2" months ("x" being the same threshold parameter as used earlier to determine the criteria for being excess).

Another factor that may occasionally be considered is the age of the inventory being used to source the order. In this case, the alternative that uses the oldest inventory should be favored. In order to compute this factor, the ages of the inventory in all the alternatives must be considered. The age is defined as the current date minus the date at which the product was added to the inventory at the location. The relative "cost" is then the maximum age minus the age of the candidate times the daily carrying cost amount for that product at that location. The following table gives an example:

Alternative	Age	Daily Carrying Cost*	Relative "Cost"
1	15	$0.0100	$0.72
2	5	$0.0140	$1.15
3	45	$0.0110	$0.46
4	87	$0.0100	$0.00

* Location Dependent

The youngest alternative (5 days) has the largest "cost" while the oldest alternative (87 days) has zero "cost". This type of factor would be used if the company were getting concerned about the product becoming excess.

Computing the Total Weighted Grade

The system will maintain a table showing the weighting factor that is to be used for each evaluation criteria. The weighted total cost to Grainger is computed and the total cost to the customer is computed. The total cost to the customer includes the extended product cost (using customer-specific prices) plus any shipping to be charged to the customer plus any special charges assess to the customer for the particular level of service the customer selects for the ADN.

Selecting the "Best" Alternative

The best alternative must meet a gross margin threshold for the preference category of the customer. Here gross margin is computed for the alternative as the total revenue received by Grainger (product price plus customer-paid shipping plus any special ADN Level of Service charges) minus the alternative's activity costs. If the gross margin threshold is not met, this means that the solution is not to be considered profitable for Grainger. The alternative is disqualified and the customer will be told that the order cannot be filled. For the ADN situation, the Grainger revenue is computed as though the customer has a 100% probability of need and will take delivery of the product. This type of check is important because some alternatives may involve so much handling and intermediate shipping that the fulfillment plan becomes too expensive and unprofitable.

If the alternative with the lowest total weighted cost to Grainger also meets the gross margin test and all of the customer's delivery and price quote requirements, then this becomes the "best" alternative.

The customer master file will have a parameter which indicates what the tolerance is between the order quoted price and the actual fulfillment cost to the customer (caused by either substituting items or having to fulfill from secondary sources when the customer pays shipping costs). If this tolerance must be exceeded, then communication with the customer will be needed to receive authorization to fulfill with the higher price levels. If no communication can be established, then the solution must be disqualified.

After some of the iterations, if there is no "best" alternative that meets the customer requirements, then the next steps depend on the ability to communicate with the customer.

If communications are not possible, then those line items which cannot be fulfilled within the customer's requirements constraints are eliminated leaving only a partial fulfillment of the order. The partial order is re-evaluated and then determined to see if it now falls within the customer's requirements.

If communications are possible, then data must be solicited regarding the ability to consider upgrade/downgrade product substitutions and/or the ability to backorder line items (or the entire order). The backorder time to ask the customer will be either the time to fulfill the problem line item from the preferred remote sourcing point (the alternative with the lowest cost to Grainger) or the time required to order from the supplier and either drop ship directly to the customer or deliver through the Grainger logistics network. (Note that the preferred remote sourcing point is also likely to have the quickest delivery time.)

Second Iteration Plan

The second iteration plan involves considering all the distribution centers not considered in the first iteration plan plus considering a special order from the supplier with either a direct ship to the customer or a delivery through the Grainger logistics network (depending on the supplier's capability to do a direct shipment). The second iteration plan does not require communications with the customer regarding substitutability or backordering. Rather, it works with the same set of customer delivery requirements as was used in the first iteration plan.

If this iteration does not produce a fulfillment plan, then communications with the customer will be required. If communications are not possible, then only those line items in the order that can be filled will be filled.

Third Iteration Plan

The third iteration plan is used if the customer is able to permit a substitution for the line item that cannot be filled in either the first or second iterations. The customer must instruct the system as to which product attributes can be varied and by how much when searching for alternatives. The system will then search for substitute products and repeat the steps for iterations one and two. If the customer has supplier diversity requirements, the system will takes these into account when identifying suitable substitutes.

If this iteration does not produce a fulfillment plan, then the fourth iteration will be considered if the customer can accept a backorder delay for either the line item in question or for the entire order.

Note that the product substitution may cause the customer cost threshold to be exceeded. In this case, the customer must be presented with this third iteration plan candidate together with the results of the fourth iteration plan. The customer can then choose to either pay the higher price for the third iteration alternative which will meet the delivery date requirements or wait for the lower price with a backordered fulfillment plan from iteration four.

Fourth Iteration Plan

The fourth iteration plan will have the delivery date requirement modified as described in the section on costing and evaluating fulfillment alternatives. This revised delivery time will have been set to either the time from the "best" alternative that exceeded the original delivery time or will be set to the time required for a special order from the supplier if there was no alternative using Grainger sources. The applicable alternative will then be designated as the fulfillment plan.

Concluding Comments

As complex as this intelligent order fulfillment plan seems, it needs to be considered that in the MRO industry, the vast majority of the orders are for single line items that will be fulfilled in either iteration one or two. Therefore, there will not be a large number of alternatives that are created and evaluated. Only a very small fraction of the orders involve more than four line items consequently resulting in a larger number of alternatives generated and evaluated.

The branch and bound structure of the four iterations will tend to limit the number of alternatives that need to be generated and considered. Even if iterations three and four are utilized, they are reached generally because there is only extremely limited or no product availability within the logistics network. Therefore the previous iterations one and two will not have created a large number of alternatives because there will not have been a large number of possible sourcing points with available inventory.

The physical implementation of the software can also be designed for parallel processing so that the generation and evaluation of alternatives can be done in parallel with each other using separate processors. This software engineering technique will reduce substantially the overall latency time needed to accomplish the intelligent order fulfillment where high performance and high transaction volumes are required.

The advantage of the graded alternatives approach to order fulfillment rather than the pure algorithmic approach is that the pure algorithmic approach does not guarantee the "best" solution. The algorithm generally stops as soon as it finds a solution. Whether that first solution is the best or not always "depends" on situations that usually not programmed within the algorithm. The graded alternative approach finds other alternatives and picks the best from the entire group of generated alternatives. Even if repetitive executions of the intelligent order fulfillment result in the same set of alternatives, the one evaluated as best may not be the same in each instance because of changes in circumstances between runs. Also, because the graded alternatives approach is generally more table and parameter driven than the pure algorithmic

approach, the company has a great deal of additional flexibility in changing, from time to time, the criteria for determining the best solutions. It should be noted that the graded alternatives approach is not unique to this application. It is a fairly common approach for searching for optimal solutions among a large universe of possibilities.

MRO ISCM – Intelligent Agent Architecture

The MRO ISCM is a comprehensive end-to-end supply chain environment that connects customers, the distributor, and suppliers together so that each has cognizance of the entire operation of the supply chain as it pertains to its own business interests. The automation of process and work is very important since the historical reputation of MRO purchasing and MRO inventory management has been that of incurring excessively high administrative and bureaucratic costs relative to the value of the underlying MRO product. Two prerequisites are needed to make this concept viable from a business success viewpoint: (1) The first is a comprehensive and reliable data communications network that serves as the foundation of connectivity between the supply chain partners; (2) The second is a collection of functions and features that make the operation and management of the supply chain as automated as possible. This paper describes both of these prerequisites.

1 The Comprehensive and Reliable Network

In general, the MRO ISCM network utilizes the Internet because of its ubiquitous accessibility and cost advantages. However, there may be situations where specific pairs of partners may choose to implement the network between them as a private network (for increased reliability and security purposes). Examples include Virtual Private Networks (VPNs) and the use of Electronic Data Interchange (EDI) networks. The architecture supports the use of both public and private networks.

The network utilizes security techniques that have become customary practices for electronic business. These include, but are not limited to, firewalls, encryption, authentication certificates, directory-based user registration and security management, etc. Because the capabilities and best practices of network and communication security are constantly evolving and improving, this document does not specify the use of any particular technique, technology or product. Instead, it simply specifies that the network architecture will support the use of security practices necessary to protect the business interests of the participants and to insure the overall integrity and confidentiality of the supply chain.

The network utilizes TCP/IP as the foundation communications protocol. Generally, HTTP and HTTPS is utilized on top of TCP/IP as the message transport envelope. These two protocols are able to deal with firewall technology better than other message management techniques. However, supply chain partners may choose to use a message-queuing system instead of HTTP and HTTPS if greater communications reliability is needed. An example of a message-queuing system is IBM's MQ-Series or the Microsoft Message Queue (MSMQ). The architecture supports the use of both HTTP/HTTPS and message-queuing systems.

All electronic data communications takes place over this network. However, for purposes of this document, the network will be divided into two logical groupings: the network for agent interaction and transaction passing, and the network used for human-based intellectual and cognitive collaboration activities. The passing of messages between agents and the passing of business transactions between supply chain partners require the same network and communication services. Therefore, the agent network is extended slightly in its specifications to embrace the communication of business transactions. Collaboration, on the other hand, involves different kinds of communication services. These are frequently more synchronous in character than those of the transaction domain. Examples of synchronous collaboration communications are the use of chat rooms, simultaneous use and maintenance of collaborative spreadsheets, blackboards, and so forth. Thus, the collaboration network has been logically differentiated from the agent/transaction network. But it should be noted that the intelligent agent domain has a tight integration with the collaboratories. Agents monitor the collaboration activity, and the collaborators have access to the services and data/knowledge of the agent domain.

2 The Intelligent Agent Network (See Diagram 1)

The intelligent agent network connects the supply chain partners so that business transactions and agent messages can be exchanged throughout the supply chain.

The network is controlled and managed by an Agent Network Manager and Broker (abbreviated hereafter simply as the "Broker"). The Broker acts as a hub. All message traffic, whether it be agent messages or business transactions, is routed through the hub. The entire network operates on the hub and spoke principle.

It should be noted that within the information systems industry, a distinction is made between mobile software agents and stationary agents. Mobile agents are capable of moving around from point to point in the network. The MRO ISCM does not utilize mobile agents. All of the intelligent software agents are of the stationary category and are located on the spokes of the network. There are two reasons for this. First, there appears to be no need for mobile agents in this business application. Second, the systems environment among the supply chain participants (customers, distributor, and suppliers) is very heterogeneous and mobile agents work best in a highly homogenous environment. For example, mobile agents are best written in Java to achieve operating system platform independence. But not all of the spoke environments support the Java environment. Furthermore, there is no homogeneous method of interfacing with the supply chain partners' legacy systems.

2.1 Agent Servers

A collection of software called an "Agent Server" is located on the spoke of this network. The Agent Server is the host for the stationary agents that perform the various pieces of functionality required for the partner. Agent Servers are used for customers, the distributor (Grainger), and suppliers. The Agent Server is also what interfaces with the partner's legacy systems. Examples of these partner[1] interfaces are:

- *Customer System Interface Examples (Left Side of Diagram 1)*
    - **Condition Monitoring** Systems which are monitoring the condition of the customer's plant equipment and facilities to determine when failure may occur and when maintenance work orders should be created and performed.
    - **Computerized Maintenance Management Systems** (CMMS) which are sometimes called **Enterprise Asset Management Systems** (EAMS) which actually manage the maintenance work.
    - **Procurement** Systems (when they are not part of the CMMS) which place purchase orders and manage the ordering process through product receipt.
    - **Accounting and Financial Settlement Systems** which account for the customer's assets and inventory, and are used to pay invoices and handle the other aspects of financial settlement of business transactions.

- *Distributor System Interface Examples (Center of Diagram 1)*
    - **Supply Chain Planning** Systems which are used to plan the overall business of the supply chain from the distributor's perspective and are used to collaborate with the planning activities done by customers and suppliers.
    - **Inventory Management** Systems and other components of the distributor's **Logistics Management** Systems.
    - **Order Management** Systems which accept and process customer orders.
    - **Accounting and Financial Settlement Systems.**
    - **Intelligent Order Fulfillment** Systems which determine the best fulfillment approach for customer orders and then manage the progress of the fulfillment activities.

[1] The terms "partner" and "participant" are used synonymously in this paper.

   - **Equipment Knowledge Base Systems** which maintain and operate the Equipment Knowledge Base required for dealing with uncertainties and probabilities of need for MRO product.
   - **Supply Chain Execution Systems** which monitor and manage the interactions over the supply chain with suppliers and customers to insure that the supply chain transactional processes are being performed correctly and timely.
   - **Supply Chain Performance Systems** which monitor and evaluate the overall performance metrics of the supply chain .

- ***Supplier System Interface Examples** (Right Side of Diagram 1)*
   - **Supply Chain Planning Systems** which are used anticipate product demand and subsequently schedule manufacturing/sourcing of product.
   - **Inventory Management Systems** which are managing the finished goods inventory.
   - **Order Management Systems** which accept orders from the distributor and processes them.
   - **Accounting and Financial Settlement Systems.**

There is also an Agent Server which is used to interface with the systems of the transportation carriers primarily to receive shipping status. (Typically, suppliers and distributors have their own dedicated systems which interface with the shipping carriers to schedule shipments, create manifests, etc.)

As noted on Diagram 1, these lists are simply examples of the kinds of legacy systems that may be interfaced to the supply chain network through the Agent Servers. The exact legacy systems will depend on the case-by-case situation of each supply chain partner.

2.2 *The Broker (Bottom of Diagram 1)*

The Broker's functions include:

- Monitor the operational status of each agent server and the agent components on that server.
- Monitor the services that each agent server and portfolio of agent components will advertise as being available from that spoke. This is used to partially determine the routing of various requests and transactions. Not all services will necessarily be available at all spokes. For operational economy purposes and other pragmatic purposes, some services (for example, the natural language translation of maintenance completion reports) may only be available on certain server spokes. These then are shared by the other spokes.
- Agents on the Agent Servers often will submit agent requests/messages and business transactions. The Broker acts as an intermediary to handle the receipt of these messages and transactions from the originating spoke, insuring that they are routed to appropriate destinations, and monitoring the status of the request fulfillment and transaction completion.
- Operate the transaction integrity services. These are sometimes referred to as the "two-phase commit" services which insure that all hubs and associated systems process their respective portions of the business transaction completely and correctly before the transaction is considered complete. These services also manage the "back out" of partially completed work on a transaction if the transaction is unable to be totally completed. This two-phase commit activity extends across all spokes involved with the completion of the business transaction.
- Operate the publish and subscribe services. Agents can both publish data and information to the network and ask to subscribe to the publishing events of certain kinds of data and information. Publish and subscribe services are typically used to monitor status and states throughout the supply chain. An agent will subscribe to a type of information publishing event if the occurrence of that event will cause the agent to take responsive or reactive action.

On the other side, the agent will publish information and data if that information and data is needed by other spokes to adjust or further process their work. (Agents determine what type of information to publish and subscribe to based on directives programmed in the AIP scripting for specific types of communications and interactions. These are discussed in a later section covering the AIP – Agent Interaction Protocol.)

- Translate between the different Knowledge Exchange Protocols. While a preference might exist to use KQML as the primary protocol, some hubs may prefer or require the use of other protocols (such as ACL, or Hive from the MIT Media Lab, or even EDI protocols for transaction traffic). By having this middleware function at the hub, there is no longer a requirement that both sender and receiver use the same protocols.
- Translate between different vocabularies. Vocabularies will be described in a later section. This service is a middleware function similar to the protocol translation function above.
- Insure and monitor reliable message delivery (both asynchronous and synchronous) throughout the network. If a particular spoke goes down, the hub will queue the message and transaction traffic for delivery when the spoke again becomes operational. If the message gets garbled or corrupted during transmission, the hub will re-transmit the message. If the message gets lost and never delivered (and this happens occasionally with HTTP and HTTPS transport protocols), the hub will re-transmit the message. Whenever re-transmission is done, the hub will insure that the message is delivered once-and-only-once. In this capacity, the hub works with its communication component counterparts at the spokes.

The Broker will utilize several key databases such as:

- The Master AIP Directory. The Agent Interface Protocol (AIP) is a specification of kinds of messages that need to be passed back and forth between partners when a major communication event type takes place. It is described in more detail in a later section of this document. A major communication event is often comprised of many task-oriented communication sub-events which represent the conversations that take place back and forth to complete the processing of the overall, major communication event. The AIP also specifies what actions and responses the senders and receivers must take for each kind of message within the major communication event. The AIP is fashioned after and is a variation of the RosettaNet Partner Interface Process (PIP). This directory consists first of a catalog of all the major communication events that are supported over the network. It also has the specific definition and specifications of content and actions associated with the sub-messages of each communications event. This will permit a spoke to download the current version of the AIP if something has happened to the local AIP database at that spoke. The Broker also uses the AIP specifications to determine what has to be monitored and managed for each communication event from the hub's perspective.
- Agents and Roles Directory. This directory lists all the Agent Servers and the components on those services. It also includes the roles or functions that the particular Agent Server plays in the network. The Broker also uses this directory to maintain the readiness status of each spoke and service on that spoke.
- Brokerage Status Database. Whenever a message or business transaction is passed through the Broker at the hub, an entry is made in this database. As fulfillment or completion progress is made by the various spokes and communicated back to the Broker, the Broker will record the changed state and status in this database. When the message of transaction is completed, the Broker will remove the entry from this database and place it on an archive audit/log database or file.
- Publish/Subscribe Database. The Broker uses this database to keep track of who the current publishers are and who the current subscribers are for each type of publishable communication.

There are other functions that the Broker performs that are not shown on Diagram 1. These include functions such as maintenance utility functions that the systems administrators use to maintain the various

databases and parameter tables, the systems management functions which the system administrators use to monitor and manage both the Broker and the overall state of the network, etc.

[note – *There is nothing particularly unique regarding the design and use of this Broker. Such brokers can often be implemented using commercially available Application Servers and frameworks.*]

3 Architecture of the Agent Server (See Diagram 2)

The Agent Server is a collection of software systems that comprises the intelligent agents located at the spokes on the network. The components can be classified into three layers: Communications, Agent Standards, and Application Tasks.

3.1 Communications Layer (The Layer at the Bottom of Diagram 2)

The communications layer is the means of connecting to the network. At the lowest level are *the basic communications protocol services*: HTTP and HTTPS will be used over the public Internet. They may also be used over the private networks although many private networks simply use TCP/IP Sockets.

Above the protocol services is the *XML parser*. All message and transaction traffic uses XML as its formatting standard. For incoming messages, the XML parser unbundles the XML formatting so that the message can be handled in its constituent parts by the knowledge exchange protocol and communicator layers. For outgoing messages, the XML parser takes the outbound constituents parts and bundles into an XML document. The XML document is then passed to the communications protocol services for transmission to the hub and ultimately to the destination spoke.

Above the XML parser is the *knowledge exchange protocol services*. These services handle the mechanics of interchange between the agents and provide the control coordination with their counterparts on the hub and other spokes. There are several knowledge exchange protocol that can be used. Examples of these are KQML, ACL, and Hive. Any one Agent Server will use only one of these protocols. Since KQML is the preferred protocol, the remainder of this document will be written as though it was sitting on top of KQML. This layer handles certain "housekeeping" chores, determines what kind of exchange "conversation" is taking place and what needs to be done for that conversation. It also unbundles the actual content of the message.

The *communicator services* is a very thin lay on the top of the overall communications layer. It logs all messages received but not completed. If completion is not acknowledged within a designated period of time, the communicator will query the higher levels of the Agent Server to determine what has gone wrong in processing the message. This layer also maintains the audit trail of completed message traffic.

3.2 Agent Standards Layer (Middle Layer on Diagram 2)

The agent standards layer has two parts: the *Domain Manager* and the *AIP Manager*. The purpose of this layer is to insure that all of the prescribed messaging and communication standards are being followed.

The *Domain Manager* takes the content and message type from the knowledge exchange protocol and determines exactly what needs to be done with the content. It first parses the message content according to the grammar and syntax of the language being used. If the message is a sub-message which is part of a conversation stream already in progress, it then passes the parsed content to the AIP Manager indicating to which AIP-in progress and what sub-message the content belongs. If, on the other hand, the message is the start of a completely new conversation, then the Domain Manager determines what server role is involved and what AIP should be activated. It then instructs the AIP Manager to start a new conversation sequence for that message and role.

The Domain Manager also responds to queries from the Broker regarding what roles and services are offered at the spoke. The Domain Manager will advertise the services and roles that are active on the Agent Server. It also monitors the operational status of the various application task agents and informs the Broker of any changes in availability of those services.

If one of the application task agents at the spoke needs to initiate communications with some other spoke in the network (for example, an ADN or purchase order needs to be sent or an order status query is being sent), then the particular application task agent will send an initiation request together with the appropriate data to the Domain Manager. The Domain Manager will determine which AIP is involved and have the AIP Manager initiate the appropriate conversation. (The AIP Manager will actually passes the message to the Communications Layer for transmittal on the network.) The same takes place when there are state change alerts at this spoke that need to be transmitted to other spokes in the network. An example of this would be the Agent Server at the transportation carriers spoke which received shipping status update information from the carrier and needs to pass this status update to the other spokes.

The *Agent Interaction Protocol (AIP) Manager* manages conversation units between the spoke and other spokes on the network. Each logical conversation unit (also called a "major communication") is usually comprised of several sub-messages that are passed back and forth between the conversing spokes. The AIP Manager directs the activity that is prescribed for the spoke and monitors the progress of the overall conversation. If a sub-message is overdue from some other spoke, the AIP Manager will initiate a query to the AIP Manager at that other spoke to determine why the delay is taking place. AIPs are described in more detail in a later section.

The AIP Manager dispatches the messages to the appropriate application task agents to perform according to what tasks are prescribed in the AIP specifications for the conversation type. When the application task agents complete their assigned tasks, the AIP Manager determines if other work needs to be done at this spoke and dispatches the appropriate tasks. If the work is done for the current sub-message, the AIP Manager associates the tasks results with the correct AIP conversation in progress and passes it to the Communications Layer for transmittal over the network to the destination spoke. If this represents the final step or sub-message for a particular AIP conversation, then the AIP Manager will inform the communicator in the Communications layer to log the overall conversation message as complete and add it to the audit trail communications log.

The AIP Manager also plans recovery from system or network outages and other systems problems that may occur either at the local site or across the network. These recovery steps can sometimes become very complex as each Agent Server must be reconciled and re-synchronized to some common state across the entire supply chain network.

Because the total supply chain system is dependent on the operations of a highly distributed network of agents and other software components that frequently utilize artificial intelligence techniques such as rules inference engines, heuristics, and other "intelligent" algorithms, conflicts can possible arise between these agents. This is because the programming and configuration of the rules files and other parametric tables and script files cannot possible cover every conceivable combination of circumstances. The overall supply chain business processes, when viewed from a holistic level, are extremely complex. When "gridlocks" and "conflicts" arise, these need to be resolved. The AIP Manager will have communication types specifically designed to negotiate solutions and resolutions to these problems and then implement the steps needed to once again get the entire supply chain functioning smoothly.

3.3 Application Tasks Layer (The Top Layer on Diagram 2)

The Application Task Layer consists of a *Task Manager* and a collection of application task agents that actually perform work.

The *Task Manager* is a rather thin layer that is responsible for insuring that non-continuous agents are started when needed and stopped when not needed, manages the system resources needed for all the agents, and generally controls the systems mechanics of what is happening when tasks are being executed. It is also the Task Manager that reports the state changes to the Domain Manager.

If one application task agent needs the services of a second agent, the request is made through the Task Manager. In this situation, the Task Manager insures that the data is correctly passed back and forth.

The individual application task agents shown on Diagram 2 can be classified according to the type of services they perform. Any given agent may, in fact, embrace multiple types of services. A typical taxonomy used by the software agent research community is:

- **Information Agents** – These agents locate information sources, extract information from those sources, provide necessary filtering of the information for relevance, and prepare the resulting information for return to the requester.
- **Integration Agents** – These agents work in the other direction from information agents. These agents receive information and data from outside the spoke and add it to the appropriate legacy system or database at the spoke.
- **Cooperation Agents** – These agents take plans created by planning agents and then direct the Task Manager to engage the appropriate agents to accomplish the task steps. If the solution involves interacting and requesting services from another spoke, the cooperation agent will determine what needs to be done, and dispatch a request to the Broker. The Broker, in turn, will figure out where the appropriate responding or cooperating spoke is and engage its services.
- **Planning Agents** – These agents develop plans and strategies. They plan out complex tasks that do not have a pre-defined execution path or AIP defined. An example of a planning agent would be at the customer site. If a shipment of a critical repair part is delayed significantly to completely ruin a maintenance schedule, the planning agent will first determine how to define an alternative sourcing for the part and what corrective actions need to be done with the maintenance scheduling and execution work. It will then have a cooperation agent initiate the corrective work. This most likely involves cooperation between several spokes on the network as the customer site needs to converse with the distributor site to determine alternative sourcing options, availability, and delivery lead times.
- **Transaction Agents** – These agents handle business transaction data, usually acting as interfaces to the site's legacy application systems.
- **Believability Agents** – These agents are basically simulators that will test certain suggested plans or hypotheses. With the MRO ISCM environment, these agents are likely to be hosted by the distributor's site/spoke, but can be accessed and utilized by other partner spokes.
- **Assistance Agents** – These agents serve as personal assistants to the humans. They help locate and retrieve information. They provide performance support assistance (context sensitive help). They provide other human interface services such as the maintenance application services to maintain the ontology files.
- **Anticipation Agents** – These agents monitor situations and anticipate developing problems.

Most application task agents will be developed and supported by the company or organization that is sponsoring and operating the supply chain system. In many cases, this will be the distributor. This will insure that there is as much consistency as possible across the entire network. However, there are situations where those agents that interface to legacy systems will be written by either the spoke company or by the software vendor supporting the legacy system.

[note – *There is nothing particularly unique about the design and use of the Agent Server. Many aspects of the Agent Server can be implemented using commercially available Enterprise Application Integration (EAI) packages and/or object frameworks such as Microsoft COM+and .NET, OMG's CORBA ORB, or Enterprise Java Beans (EJB).*

What is unique in this design is the close coupling of the Agent/Transaction Network with the Collaboration Network such that the agents can enlist the services of the collaboration's workflow management system if needed and the collaboration participants can access ontology knowledge and use certain services from the Agent Server. This is depicted in the upper right corner of Diagram 2 which shows the Collaborative Interface agent.

Another area of unique design is the extension of the agent architecture to embrace the business transactions within the same logical network. This enables the agents to work closely with the actual business transactions. It also allows a single common Server and Network architecture to service both types of message traffic thereby eliminating the need for rather redundant systems software components to run the servers and network.]

3.3.1 Application Task Agents Likely to be Found at Customer Spokes

The diagram uses a customer spoke context. At the left are a series of *legacy interface agents*. There will be one agent for each legacy system that needs to be integrated into the supply chain network. The ability to integrate a legacy system and the permission to extract and add information to that legacy system need to be expressly stated in the legal agreement that establishes participation by the customer in the supply chain system. These agents are classified as Information, Integration, and Transaction Agents.

The *user interface agents* will be servlets, scripts, or CGI programs running on a web server. These enable users at the customer site to interact with the agent environment using web technology and browsers. Such interface is necessary for users to determine the status of certain events and transactions. It is also necessary when parts of the ontology need to be changed and otherwise maintained. Local systems managers will use this interface to query the Domain Manager to determine the operational status of the various component of the Agent Server. These agents are classified as Assistance Agents.

The *algorithm executor agents* perform all kinds of analysis, synthesis, and evaluation. These are classified as Planning Agents, Cooperation Agents, Believability Agents, and/or Anticipation Agents. Assistance Agents that are performance support routines will also be algorithm executors.

The *rules inference engine* is more of a functional component rather than a software agent. The rules inference engine is used whenever decisions need to be made that follow rules and policies documented in rules files contained in the ontology. All of the other agents can both use the services of the rules inference engine and will also likely be asked to contribute data to queries by the engine.

The *knowledge manager* maintains the ontology. Again, this is not so much an agent as it is a functional component of the Agent Server. The ontology is a critical component of the agent environment since all of the knowledge that the agents act upon is kept and cataloged in the ontology. All of the other task agents will likely both query and update the ontology using the functional services of the knowledge manager.

The *collaboratory interface* consists of several types of agents. If other task agents need to involve a user, the collaboratory interface is used to make entries into and get results out of the workflow management system. For example, a revised sourcing and fulfillment plan might cause the price of the repair parts to exceed an approval threshold and require that an entry be place in the workflow system to have the new plan approved. The appropriate information is packaged as a business process event and submitted to the workflow management system. The workflow management system routes the package to the appropriate people for approval or disapproval and returns the decisions to the collaboratory interface. The interface will pass the decision results to the agent to continue processing. Another purpose of the collaboratory interface is to enable the actual software tools used in collaboration to interact with the Agent Server. The collaborators may want to query areas of the ontology, or they may want to use Believability Agents to run

a simulation. Finally, the collaboratory interface allows an Anticipation Agent to monitor the well-being or health of the collaboration itself by monitoring and analyzing collaboration activity taking place.[2]

3.3.2 Application Tasks Likely to be Found at Supplier Spokes

The portfolio of application tasks needed by a supplier is much smaller. There will be legacy interface agents for the supplier's supply chain planning/forecasting systems, the inventory management systems, the order management systems, and the accounting/finance systems. Suppliers may request a limited user interface web server to give them visibility into the status and operation of the Agent Server. Algorithm executor agents and the rules inference engine will only be used by those suppliers who want to develop a very sophisticated involvement in the supply chain operations. The ontology for the supplier site will be much more limited than that at the customer site. Whether or not the supplier has a collaboratory interface will depend on the degree to which the supplier wishes to truly collaborate in the supply chain.

3.3.3 Application Tasks Likely to be Found at the Transportation Spoke

The transportation spoke essentially requires only the legacy interfaces to the various carriers' systems, a small ontology to record which shipments should be tracked for status reporting, and a cooperation agent to periodically query the shippers' systems to obtain the status and publish the status on the agent network so that the distributor and customer spokes can be advised of the status. There may be a basic anticipation agent executor that uses the rules inference engine to anticipate major transportation problems that could be developing based on the status that is being currently received.

3.3.4 Application Tasks Likely to be Found at the Distributor Spoke

This is clearly the largest collection of application tasks. All of the same kinds of application task agents found at the customer site will be found at the distributor spoke: legacy interfaces, user interface web server, algorithm executors, the rules inference engine, a robust ontology, and a collaboratory interface. Types of legacy systems that require interfacing include supply chain planning, inventory management, order management, accounting/finance, intelligent order fulfillment, the equipment knowledge base, the supply chain execution manager, and the supply chain performance system. There will be more algorithm execution agents because the distributor will have much more to accomplish in coordinating the happenings between the supplier and customer.

4 Knowledge Exchange Protocol – KQML

[note – *Clearly there is nothing unique in the use of knowledge exchange protocols such as KQML or ACL. What is unique in this application is the extension to embrace business transaction events in the same agent communication framework. This enables both the business transactions and the agent conversations to reference each other. It is the basic means by which the agent community can become involved with the actual business transactions and transaction processing events.*

Another point of unique design is the addition of the ":aip:" performative parameter. The tying together of KQML specifications and AIP execution specifications provides a much more comprehensive and consistent specification of how the agents should perform and how the business transactions are to be processed and integrated with the agent activity.]

[2] There is a separate white paper describing collaboratories and methods to monitor the health of the collaborative affiliations between supply chain partners.

As mentioned earlier, the preferred knowledge exchange protocol is KQML developed in the early 1990's at the University of Maryland. The full name of the protocol is Knowledge and Query Manipulation Language. The MRO ISCM system actually uses our own adaptation and extension of KQML, not the true and exact specification.

KQML defines both a message format and a message handling system for multi-agent systems. The message transmission definitions are only virtual. Instead of defining a particular technical standard for implementing message transport, KQML only defines a model of the services that the message transport system must perform. Basically, the message transport must provide point-to-point, reliable messaging. However, the pure KQML model was modified to meet the needs of incorporating both agent messages and business transactions in the network environment. The KQML message transmission definitions are incorporated and modified as needed in the Communications Layer of the Agent Server. KQML covers identification, connection establishment, and message exchange. However, the semantic content of the message is not covered. KQML is more of an envelope for the semantic content. The semantic content can be expressed in any knowledge language. Examples of knowledge languages are the Knowledge Interchange Format (KIF), Conceptual Knowledge Markup Language (CKML), traditional programming languages such as Lisp and PROLOG, declarative languages such as SQL, etc. Essentially, any kind of language can be used. The requirements are: a) that the language representation be sentential – expressions using the representation can be viewed as entries in a knowledge base; and b) that sentences have an encoding as an ASCII string so that they can be embedded in the KQML message.

KQML is built around a set of performatives. A performative is a fundamental action that one agent can request another agent to perform. The performatives are broken into major groups:

- Basic Informative Performatives
  - Tell – the content sentence is in the sender's knowledge base
  - Deny – the embedded performative no longer applies to the receiver
  - Untell – a deny of a tell
- Database Performatives
  - Insert – the receiver should add the content sentence to its knowledge base
  - Delete – the receiver should delete the content sentence from its knowledge base
  - Delete-One – the receiver should delete one sentence from its knowledge base that matches the content sentence
  - Delete-All – the receiver should delete all instances from its knowledge base that match the content sentence
- Basic Responses
  - Error – the sender cannot understand the designated performative which it previously received or the performative contained errors that made it illegal for the sender to execute
  - Sorry – the sender understands but is not able to provide a response to a performative
- Basic Queries
  - Evaluate – the sender would like the recipient to simplify the expression in the content parameter and reply with the result
  - Reply – this is the response that the sender believes is the correct reply to a Evaluate performative it earlier received.
  - Ask-If – the sender wishes to know if the recipient has any matches in its knowledge base that correspond to the schema in the content parameter
  - Ask-About – similar to the Ask-If except that the response is a collection of all the sentences in the receiver's knowledge base that match the schema
  - Ask-One – similar to Ask-If except that the reply should be formatted to a different schema than the schema in the query

   - Ask-All – similar to the Ask-One except that the reply is a collection of all the sentences that match the query schema but reformatted to a different schema
   - Sorry – same as the Basic Response
- Multi-Response Queries
   - Stream-About – similar to an Ask-About except that the reply is a collection of matches that contain a series of performatives that when taken together identify the members of the collection
   - Stream-All – similar to the Ask-All except that the reply is a collection of performatives that when taken together identify the members of the collection
   - EOS – "End of Stream" which indicates the end of the series of performatives that respond to a Stream-About or a Stream-All request.
- Basic Effector Performatives
   - Achieve – a request that the recipient try to make the content sentence true in the context of its knowledge base
   - Unachieve – a deny of an Achieve
- Generator Performatives
   - Standby – the receiver should announce its readiness to receive further performatives
   - Ready – a response to a Standby performative
   - Next - the sender wishes to receive the next response
   - Rest – the sender wishes to receive the remaining responses in a stream from those promised by a Ready performative
   - Generator – a combination of Standby for a Stream-All
- Capability-Definition Performatives
   - Advertise – indicates that the sender is particularly suited to process a class of performatives
- Notification Performatives
   - Subscribe – the sender wishes the recipient to tell it about future changes that would be a response(s) to the performative in the content parameter
   - Monitor – similar to a Subscribe except that the response is a Stream-All
- Networking Performatives
   - Register – indicates the sender can deliver performatives
   - Unregister – a deny of Register
   - Forward – the sender is to execute the content performative but send the response to a different agent
   - Broadcast – the recipient should route the results of a performative to a list of agents
   - Pipe – indicates that future traffic on this channel should be routed to a different agent
   - Break – breaks a pipe
   - Transport-Address – the recipient should use a specific message transport address rather than the agent's identification for return responses
- Facilitation Performatives
   - Broker-One – the sender should process the performative through the help of a single agent that is particularly suited to processing the embedded performative
   - Broker-All – same as Broker-One except that a list of candidate helpers is provided
   - Recommend-One – the recipient should respond with the name of a single agent that is particularly well suited to the processing of a performative
   - Recommend-All – same as Recommend-One except that the response is a list of all candidate agents

- Recruit-One – the recipient should forward the request to a single agent that is particularly well suited to the processing of a performative
- Recruit-All – same as Recruit-One except that the performative is forwarded to all qualified agents

The definition of performatives is open-ended in KQML. Implementers can add specialized new performatives as long as the parameters and syntax comply with the KQML specifications. The MRO ISCM system requires that an addition group of performatives be added that can be called "Business Transaction Performatives". This new group allows business transactions to be handled over the agent network as well as agent message traffic and also allows the agents to interface with and become involved with the business transactions.

- Trigger-Event – the recipient is receiving a business transaction that is should process but is not subject to transaction semantics (two-phase commit)
- Start-Transaction – the recipient is receiving a business transaction that it should process and the transaction is subject to transaction semantics
- Add-Event – the recipient is receiving another transaction in the set or series embraced in the currently active Start-Transaction
- Commit-Transaction – All agents have successfully processes the transaction events in the logical Start-Transaction set and can now commit and all agents can now commit their results
- Abort-Transaction – the sender is unable to successfully process a transaction event in the current Start-Transaction series and therefore all agents participating in the currently active Start-Transaction set should abort their efforts and back-out any results already completed

Each performative has a series of parameters identified by keywords. Some of these parameters are (this is not an all-inclusive list):

- :content<sentence or embedded performative> this is the "direct object" of the performative title
- :sender<name> this identifies the sender of the performative
- :receiver<name> this identifies the recipient of the performative
- :language<text> this is the language in which the :content is expressed
- :ontology<text> this identifies the ontological subject that the recipient should involve within its ontology database when executing the performative
- :force<word> this indicates that the sender guarantees that it will never deny the meaning of the performative if force is equal to "permanent". If no such guarantee exists, then force will be set to "tentative".

Here is an example of a performative that the distributor may send to a customer to inquire about the age of a component in the equipment being maintained and the subject of a pending ADN:

```
(Evaluate
        :content(Age of part xyz in equipment abc reference work order qqq)
        :language(English)
        :ontology(Equipment)
        :reply-with(reference36972)
        :sender(Grainger-ReliabilityCenter)
        :receiver(Customer-A)
)
```

The customer's spoke would reply with this performative:

```
(Reply
        :content(4588 production cycles)
        :language(English)
        :ontology(Equipment)
        :in-reply-to(reference36972)
        :force(tentative)
        :sender(Customer-A)
        :receiver(Grainger-ReliabilityCenter)
)
```

When a customer submits a purchase order the performative would be:

```
(Trigger-Event
        :content(an XML document with the purchase order data)
        :language(XML)
        :ontology(Orders)
        :reply-with(PO77983AX2)
        :force(permanent)
        :sender(Customer-A)
        :receiver(Grainger-OrderEntry)
)
```

Grainger would send a confirmation with this performative:

```
(Reply
        :content(an XML document with the order confirmation number and info)
        :language(XML)
        :ontology(Orders)
        :in-reply-to(PO77983AX2)
        :force(permanent)
        :sender(Grainger-OrderEntry)
        :receiver(Customer-A)
)
```

The customer would immediately want to subscribe to any order status messages:

```
(Subscribe
        :content(Reply
                        :content(name of XML DTD/Schema for order status*)
                        :language(XML)
                        :ontology(Orders)
                        :in-reply-to(PO77983AX2)
                        :force(?)
                        :sender(?)
                        :receiver(Customer-A)
                )
        :ontology(Orders)
        :language(KQML)
        :reply-with(PO77983AX2)
        :force(tentative)
        :sender(Customer-A)
        :receiver(NetworkHubBroker)
)
```

* The XML DTD/Schema contains actual default data element values that identify the particular order to be monitored. This data will also be returned in the XML order status document.

In the last example, the NetworkHubBroker will, in turn, send out performatives to the Grainger-OrderEntry spoke to find out the status and subscribe to status change events. When shipping reference data is known, the NetworkHubBroker will also send out performatives to the Transportation Agent spoke to monitor and publish shipping status changes. Customer-A will subsequent receive Reply performatives from both Grainger-OrderEntry and the TransportationAgent spoke.

Finally, a new performative parameter is added to the standard parameters specified in KQML. This is the ":aip<word>" parameter and is used to specify the name of the Agent Interaction Protocol (AIP) that specifies the steps to follow in executing the performative. This is described in more detail in a following section. This new parameter has not been included in the above examples because these preceed the section in this document describing the AIPs.

5 The Ontology (See Diagram 3)

[note – *The basic construction of an ontology with a semantic network, the substitution of Frames for otherwise simple concept nodes, and the inclusion of rule sets are generally understood practices in the area of artificial intelligence. What is unique with the MRO ISCM application of an ontology is the inclusion of the AIP to direct the usage and interpretation of the ontology contents. Usually, usage an interpretation is embedded in the programming code of the actual agent software. By abstracting it out of the programming code and putting it into a parameter-table-like or scripting AIP module, greater consistency in semantic interpretation across the entire community of agents is achievable. Also, usage and interpretation re-use is enabled better with the AIP approach rather than the embedding into programming code. As will be discussed in the section on the AIP, the core AIPs are defined by the supply chain sponsor and replicated consistently across the entire supply chain network. While an individual Agent Server can add its own local AIPs to facilitate unique local processing and functional requirements, any AIP that involves interaction with other Agent Servers across the supply chain network requires the use of the standard supply-chain-wide AIPs This mandates consistent semantic interpretation of the ontology contents.*]

An ontology is a description of concepts and relationships that can exist for a community of agents. The ontology exists for the purpose of knowledge sharing and knowledge re-use.

The first portion of this section provides an introductory tutorial on ontologies and is for the benefit of those readers who are less experienced with artificial intelligence and knowledge engineering.

Each Agent Server spoke will have its own ontology. The ontology is physically implemented using a database management system, so there will be a series of programs to drive the maintenance and querying of the ontology database. These are known as the Knowledge Manager in the Agent Server.

The ontology represents the concepts, relationships, and knowledge with declarative formalism. It is shown on the left side of Diagram 3. The basic objects of the ontology representing the concepts are Frames. A Frame is an object-oriented analysis/design technique that documents an object together with its attributes and operations. Many ontologies simply use plain nodes in their semantic networks. By using Frames as the nodes in the semantic network, attributes in a structured data context can be included in the ontology rather than having to represent each attribute as another node-relationship construct. For example, maintenance work orders have many attributes that are more convenient to model and use in the Frame construct than in the node-relationship construct.

The Frame construct also allows operations to be associated with the node concepts. These operations will be used for several purposes. First, some aspects of the ontology database will actually replicate data in the databases of the legacy systems. In these situations, rather than duplicate the data in the ontology, an

operation will contain the SQL and other access methods/procedures to access and maintain this data in the legacy system database. An agent may query a part of the ontology and think it is getting the data directly from the ontology when, in reality, the ontology database is simply pointing (through operations in that concept node's Frame) to the real database in the legacy system domain. The operations in the Frame are also used to direct the employment of other Application Task Agents. In this way, the ontology not only records relationships between concepts, but actually records processes that the concepts themselves use or are used by other applications in better understanding the concepts. They can also server as "database triggers" to cause certain agent and business transaction events to be triggered whenever certain types of ontology maintenance takes place against the concept.

Another purpose of using Frames is to permit a hierarchy of concepts where child concepts can inherit the attributes and operations of the parent concept. The ability to provide inheritance eliminates the need for substantial amounts of redundancy in the physical implementation of the ontology database.

Frames can represent just about any kind of concept. For this reason, the ontology will be logically (and sometimes physically) divided into partitions that represent particular subject areas. It is these subject areas that are referenced by the ":ontology" parameter in the KQML performatives. There is a base set of subject areas that is defined by the sponsor/operator of the supply chain (usually the distributor) that will be common across the entire supply chain. This will enable a base level of semantic consistency among the community of agents in the supply chain. However, each Agent Server can add its own subject areas as the need arises. The artificial intelligence community has established techniques for one ontology server to describe these new subjects in its semantic network to other agent servers so that some degree of automatic discovery can take place thereby eliminating confusion caused when one set of agents does not understand the semantic implications of a foreign ontology being used by cooperating agents.

Frames are related to each other through "relationships". Every relationship will have a dominant Frame and a subordinate Frame. The relationship will have a descriptive name for each direction. For example, in an inheritance relationship, the dominant Frame is the parent and the subordinate Frame is the child concept. The relationship is named "Has-A-Kind-Of" in the direction of dominant to subordinate and the name of "Is-A" in the direction of subordinate to dominant. So a Parent has a kind of concept for a child and the child is a parent. The "Is-A" / "Has-A-Kind-Of" relationship is the only relationship that enables inheritance between the related concept Frames.

If a Frame is composed of component sub-Frames, then the Part-Of relationship is used. "Part-Of" is the name from the subordinate to the dominant and "Has-Component-Of" is the name in the direction of dominant to subordinate.

Sometimes a dominant Frame represents an entire class of concepts and the dominant Frame represents a single instance of that class. In this case the "Instance-Of" / "Has-Instance-Of" relationship is used.

Just about anything goes for relationships. Relationships do not have to posses unique names. However, any pair of Frame types can only have one relationship with a given set of names. So, relationship names are only unique with specific pairs of Frame types.

The semantic network in the ontology also includes rules sets that are used by the Rules Inference Engine. These rule sets are collections of If-Then-Else rules used to control decision making and to enforce business policies. Rule sets also have relationships with concept frames which basically indicate that the rules pertain to the concept embodied in the Frame. Generally, some operation in a concept Frame will trigger the running of an algorithm executor agent. This agent, in turn, may need to do some decision making that requires using the Rules Inference Engine. The algorithm executor agent will retrieve the applicable rule set from the ontology by finding the rule set related to the target concept Frame. The executor agent will direct the Rules Inference Engine to retrieve the rule set from the ontology, perform the inference, and modify the ontology based on the results of the inference. The executor will subsequently learn about the outcome of the inference (that is the outcome of the decision making task) by querying the ontology. Based on the outcome of this query, the executor will branch its processing logic accordingly.

"Knowledge" and "Understanding" are achieved by the agents by traversing the ontology from concept Frame to concept Frame over the relationship connections. For example, in order to answer a query "Will the widget on equipment XYZ fail in the next 4 weeks?" an information agent will respond to an evaluation performative. The AIP will first tell the information agent to attempt to find the concept Frame for equipment XYZ and its Part-Of component Frame for the widget. It will then navigate around those two target frames to determine if there are any relationship paths to other concept Frames that indicate possible causes of wear/tear and failure of the widget. It will then navigate around the target Frames to search out any relationship paths to concept Frames that indicate the presence or absence of these potential causes. If potential causes are evident, then further relationship navigation is done to determine the severity or potential likelihood of a failure. Eventually, the evaluation agent will compile enough "knowledge" and "understanding" from the ontology to reply with a Yes or No and provide some sort of confidence metric in that reply. If no answer can be determined from the ontology, then the information agent can generate sorry performative reply.

Like any other database, the ontology will have a meta-knowledge section which describes the schema of Frames, attributes, operations, relationships, and rule sets and the permitted pairings of these. The meta-knowledge is shown on the right side of Diagram 3. The supply chain sponsor/operator (usually the distributor) will provide a core set of meta-knowledge so that there is a consistent semantic net foundation across the entire supply chain. Besides populating the ontology with specific instances of the meta-knowledge components, the Agent Server can also extend the ontology schema by making local additions to the meta-knowledge and then populating instances of those additions.

The supply chain sponsor will also provide a core set of AIPs which, among other execution directives, will define what kinds of relationship names to use in different performative actions and what kinds of relationship names to create when creating new relationships. Therefore the real semantic definition of relationships is inherent in the AIPs since the AIPs are determining how the concept Frames and relationships are being used and interpreted by the various application task agents.

6 Agent Interaction Protocols

The Agent Interaction Protocol (AIP) is patterned after the RosettaNet Partner Interface Process™ (PIP™). However, this is not a true implementation of RosettaNet. Instead, the AIP is a more detailed and expanded concept specifically designed to coordinate the execution of various agent tasks within a single Agent Server as well as coordinating the communications between agents across the entire supply chain network.

A quote from the RosettaNet document entitled "Understanding a PIP Blueprint – Release 3.1" concisely describes a PIP:

> A PIP depicts the activities, decisions and Partner Role interactions that fulfill a business transaction between two partners in the supply chain. Each partner participating in the PIP must fulfill the obligations specified in the PIP. If any partner fails to perform a service as specified in the PIP implementation guide then the business transaction is null and void.
>
> A PIP must:
>
> - have a measurable business outcome or output;
> - not contain proprietary business processes;
> - preferably contain more than one role interaction; and
> - be a discrete unit of work that can be attached and built into other PIPs to achieve a larger business outcome.

Among other things, a PIP contains business process flow diagrams, a definition of the start and end states that describe the conditions that need to exist to begin and complete the business process, the specification of partner roles, a detailed statement of business activities initiated by the partner role, a table of security,

audit and process performance controls, and a specification of the business documents both input and output from the business process.

Typically, the major business process involves multiple communications between the partners as each set of role activities progresses. Some of these communications are simply to provide acknowledgment of receipt of earlier communications. Other communications actually transmit information.

RosettaNet has defined PIPs for the major business transactional processes involved in the manufacturing industry. Many of these have direct counterparts to the business transactions carried out on the MRO ISCM supply chain.

Whereas the RosettaNet PIP embraces only business transactions, the MRO ISCM Agent Interaction Protocol is extended to embrace agent interaction as well as business transactions. The AIP defines major interaction roles that the entire Agent Server at the target spoke has. These correspond to the PIP business partner roles. The AIP specifies overall sets of activities that each Agent Server/spoke must execute before the interaction (either agent interaction or business transaction processing) can be considered complete. However, unlike the PIP, the AIP specifies a set of sub-roles that are local to the Agent Server. Each sub-role is assigned to be the responsibility of one or more of the application task agents. There are start and end states defined for each of these sub-roles. There are task activities specified that each responsible task agent must execute to complete the sub-role. There is also a specification of the input and output message content. There is a set of task activity performance controls for the sub-roles that are comparable to the business activity performance controls in the PIP. Where the PIP only covers interactions between supply chain partners, the AIP covers local interactions between the Activity Task Agents within an Agent Server as well as the interactions between different Agent Servers.

The definition of the task activities in the AIP can, at times, be very granular and detailed. The various task agents are provided with instructions and processes on how to complete their assigned responsibilities. This also includes directives regarding how to use and update the ontology. As mentioned in the section on the ontologies, the AIP actually encapsulates the actual semantic interpretation of the ontology because it specifies how the agents are to use and update the ontology. This includes how to create and assign relationship names between concept frames. It provides direction to the agents as to what concepts and relationships to look for in the ontology and what rule sets to use in making decisions.

Every performative that comes across the supply chain network will have a particular AIP type assigned to it. This insures that all the participating Agent Servers are executing coordinated and synergized roles to complete the performative. This annotation on the performative will include the type of AIP, the actual AIP instance identification, and particular sub-role and activity step for which the performative is the input or output message. This allows the receiving Agent Server to immediately pick up the processing of the task activities for the AIP in progress at the correct place in the entire AIP execution script when a performative is received.

Another distinguishing feature of the AIP is that there are circumstances when it is not possible to pre-determine exactly how many exchanges of intermediate conversations between the various agents will take place. Sometimes there are many iterations of intermediate communications until some conclusion is reached between the cooperating agents. For this situation, the AIP has an activity definition construct that specifies an iterative loop to take place until some condition is met (equivalent to the programming constructs of "WHILE (some condition) DO (this)" or "DO (this) UNTIL (some condition)". In fact, the entire specification of what task activities to perform is implemented in the AIP using a pseudo-programming language similar to the scripting languages used on Web servers.

Technicians at the supply chain partner spoke site can maintain the various pieces of the AIP including the specification of task activities by using the User Interface Web Server in the Application Task layer of the Agent Server. This can be done in order to support site-unique processing steps and policies.

7 Examples of End-to-End Flows Across the Supply Chain

The following example traces the flow of activity throughout the supply chain network from the time a customer determines that maintenance will be required on a piece of equipment until the maintenance is completed and the distributor has invoiced the customer for the MRO products used in the maintenance work.

Diagram 1 can be used visualize the flow of activity across the entire supply chain network. Diagram 2 can be used to visualize the flow of activity within a single Agent Server at a spoke. The process descriptions in this example are at the granular level of detail as used in Diagram 2. The overall flow in this example is also cross-referenced to the SCOR Model Rendering of the MRO ISCM.[3] This is an overall business process diagram covering the entire supply chain and using the Supply Chain Operation Reference (SCOR) model from the Supply Chain Council. This cross-referencing appears as the Process Identifiers in the square brackets. SCOR processes starting with the letter "C" are for the customer and appear on the left side of the diagram. Processes starting with the letter "G" are for the distributor, Grainger, and appear in the center of the diagram. Processes starting with the letter "S" are for the supplier and appear on the right side of the diagram. Processes starting with the letter "T" are for the transportation carrier and appear toward the left side between the customer and Grainger. This example does not include any involvement of a third party financial settlement institution.

1. The customer's condition monitoring system detects a deterioration of the widget on equipment XYZ and determines that maintenance should be scheduled. [CM4] The condition monitoring system interacts with the CMMS system so that the CMMS system schedules a maintenance task to be done in 8 days. [CM5]
2. The legacy interface agent to the CMMS system detects that a new maintenance task has been scheduled. It immediately triggers an internal AIP to initiate tracking of this new maintenance task. This is accomplished by having the legacy interface agent to the CMMS send an alert message to the Domain Manager. Based on the message type and contents, the Domain Manager determined the AIP that applies to the need to start tracking this maintenance task.
3. The AIP Manager opens a new AIP instance and begins assigning the activity tasks to the various agents to begin tracking this maintenance task. The end result of this set of activities is that the ontology is updated regarding the equipment, the condition of the widget, and the fact that a maintenance task has been scheduled for a certain date. The last activity in this AIP is to initiate an Evaluate performative to a Planning algorithm agent to determine how the repair parts will be sourced and assigned or reserved to the maintenance task. This is actually accomplished by having the AIP Manager initiate an AIP instance for the AIP that determines pre-maintenance planning parts sourcing. [C-P2.1]
4. The AIP for pre-maintenance planning parts sourcing directs a Planning agent to collaborate with the CMMS interface agent to determine the sourcing plan. The Planning agent determines from the CMMS system that in addition to the widget the left flapper fud will also have to be replaced if the widget in fact fails. [C-P2.2] There is a spare widget in the local storeroom, the there is no left flapper fud in the local storeroom. [CM3 and C-P2.3] The planning process also determines that if the spare widget is needed then it will be replenished from the distributor. During the planning process, several rule sets were extracted from the ontology and run through the inference engine to arrive at the final decisions and plan. Throughout the planning activity, additional detail is added to the ontology. (Recall that some of the applicable ontology contact actually resides in the condition monitoring system database and the CMMS database. Pointers to these data are actually added to the ontology.) The last step is to communicate the plan to the CMMS system so it can create the ADN. [C-P2.4]
5. The CMMS creates the ADN for the maintenance task. [C-S1.1] The ADN has two line items. The line item for the left flapper fud specifies a level of service to have the left flapper fud staged at the local branch 2 days before the maintenance task is scheduled to be done. The other line item specifies that the widget be staged for non-urgent delivery to the customer. Since this will be replacing a widget

[3] Visio file: MROISCM_SCOR.vsd::TopLevel

currently in the local storeroom, so there is no hurry in having the replacement delivered to the customer. The CMMS system requests the Domain Manger (through the legacy system agent) to initiate a business transaction performative to pass the ADN to the distributor. [C-P2.4 and C-S1.1]

6. The Domain Manager determines that the "Submit ADN" AIP will be executed and instructs the AIP Manager to start an instance of that AIP. This AIP requires that the CMMS pass the ADN transaction to a Information type algorithm executor agent which will package the ADN into the appropriate Trigger-Event performative. This performative is sent to the AIP Manager which, in turn, sends it to the Communications layer where it is sent to the distributor. However, the AIP indicates that the distributor must return an acknowledgment, so the AIP remains open.
7. The Broker sends the ADN transaction to the distributor's Agent Server. The distributor's Domain Manager determines that this Trigger-Event performative involves a new ADN with the Submit ADN AIP. The distributor's domain manager instructs the distributor's AIP Manager which opens up a new instance of this AIP. However, the AIP Manager is only interested in the roles that the distributor must play.
8. The distributor's AIP Manager instructs the order management system (through the legacy interface agent) to accept the ADN. [G-D1.1]The order management system accepts the ADN and assigns a receipt confirmation number and passes this number to the AIP Manager. [G-D1.2] The AIP Manager has an Information type algorithm executor agent package the confirmation number in a Reply performative. The Reply performative is returned to the AIP Manager which sends it to the Communications layer for delivery to the customer. All of this activity has caused certain entries to be made in the distributor's ontology regarding the existence of the new ADN, the equipment involved, and the two parts cited on the ADN.
9. The distributor AIP Manager determines from the last step of the activities for its role in the Submit ADN AIP is to start a Intelligent Fulfillment Planning AIP for the ADN. This will be continued in step 12.
10. The Reply performative containing the distributor's ADN confirmation number is received at the customer's Agent Server. It gets passed to the Domain Manager which, in turn, passes it to the AIP Manager citing the open AIP and the step that is awaiting confirmation. The AIP Manager has the Knowledge Manager add the confirmation information to the ontology. The last step of the customer's role in the Submit ADN AIP is to send out a Subscribe performative to listen for status updates on the fulfillment of the ADN. [C-EP.2] The AIP Manager has the Information type algorithm executor create the necessary Subscribe performative, and sends it to the Communications layer which then dispatches the message to the Broker. This completes the Submit ADN AIP for the customer, so it is closed.
11. The Broker receives the Subscribe performative and sets up an entry in the publish/subscribe database to service this request. The Broker also sends a Subscribe performative to the distributor to inform the distributor to send status change messages when they occur.
12. The distributor's AIP Manager begins to execute the activity tasks for the Intelligent Fulfillment Planning AIP. [G-P4.1, G-P4.2, G-P4.3, and G-P4.4] A number of rule sets will be involved in this planning activity. One of the rules requires the age of the existing widget. The distributor's Planning agent which is working with the Intelligent Order Fulfillment Planning system detects that the customer did not provide any probability of need on the ADN nor was there any age of the current widget indicated on the ADN. The activity script in the AIP indicates that under this condition, additional information about the age of the subject part must be solicited from the customer. [G-P4.1] So the Planning agent instructs the Domain Manager to initiate an Equipment Age/Condition Query AIP to solicit some additional information from the customer's ontology. This AIP specifies that an Evaluate performative must be sent to the customer. This is done. In the meantime, the Planning agent suspends any further work on the fulfillment plan until a Reply performative is received from the customer.
13. The customer receives the Evaluate performative, opens up a new Equipment Age/Condition Query AIP instance. The agent complex determines that it does not know the exact age of the widget in question. However, a rule set in the ontology enables the agents to estimate the age with a certain confidence level (albeit somewhat low!). Furthermore, the condition monitoring database has more details on the actual monitored condition of the widget. All of this information gets packaged up in a Reply performative according to the activity script in the AIP. When the Reply is dispatched back to

the distribution, the customer's AIP Manager closes out the Equipment Age/Condition Query AIP instance.

14. The distributor receives the Reply performative. Work on the fulfillment plan is resumed. The rule sets in the ontology were able to synthesize data from the Equipment Knowledge Base together with the Reply data received from the customer to guestimate that the probability of need will be set at 87%. This enables a fulfillment plan to be created. Upon completion of the Intelligent Fulfillment Planning AIP, the distributor's AIP Manager initiates an ADN Staging AIP to execute against the fulfillment plan. [G-P4.4 and G-D1.3]
15. Because the distributor's ontology indicates the customer has subscribed to status updates, ADN Status Update AIPs are initiated for two status messages for the customer. The first passes the information that the probability of need has been guestimated to be 87%. The next status indicates that a fulfillment plan has been designed and will meet the customer's level of service specifications for both line items. Since the Subscribe performative requested that the status updates be sent as Insert performatives, the distributor formats the status as Insert performatives. [C-EP.3]
16. We skip ahead to the point where a status message is sent to the customer indicating that both the widget and the left flapper fud have been staged according to the level of service specifications. Like all the other status messages, this status is added to the customer's ontology. [G-D1.3]
17. On the day before the maintenance is to be performed, the CMMS extracts all the maintenance work orders for the next day. This action is detected by the legacy interface agent which sends an alert to the customer's Domain Manager. The Domain Manager schedules a Verify Maintenance Readiness AIP. This AIP scans the ontology base for the status of the ADN fulfillment. Since the status indicates the products have been staged, the agents determine there is no need to delay the maintenance. In this case the customer has adjusted the AIP scripts not to take any action. If, however, there was a staging problem, the agents would have followed the AIP script to create a workflow event to notify the maintenance manager of the staging problem.
18. The maintenance worker begins the inspection portion of the maintenance task on the designated day. [CM6] The inspection reveals that the widget is, indeed, failing and needs to be replaced. As was predicted, the left flapper fud also will need to be replaced. The maintenance worker makes these entries in the CMMS system and notes that the widget is available from the local storeroom. So the widget is immediately requisitioned from the storeroom. However, the left flapper fud will need to be picked up at the local distributor branch where it has been staged.
19. The CMMS system immediately creates the purchase order to take delivery of the parts reserved under the ADN. [C-S1.1] The CMMS system, through the legacy agent, requests the Domain Manager to create the correct performative to immediately send the purchase order to the distributor. The Domain Manager determines that the Deliver Reserved ADN Product AIP will be used. Without going into the details, this AIP creates a high priority message containing a Begin-Transaction performative for the purchase order.
20. The distributor receives the Begin-Transaction performative with the purchase order and the designation of a Deliver Reserved ADN Product AIP. [G-D1.2] Again, without going into the details, this AIP directs the purchase order information be entered into the distributor's order management system as a high priority transaction. Also, because this purchase order is under transaction semantics control (cause by the Begin-Transaction performative), it will not be committed until the branch confirms that the left flapper fud has been place in the will-call bin and the distribution center confirms that it has added the picking and shipping of the widget to its work list. [G-D1.3 through G-D1.7 for the widget and G-D1.9 for the left flapper fud] When both of these are completed, the purchase order transaction is committed and the customer's Agent Server so notified. If there had been a problem either at the branch or at the distribution center, the transaction would have been aborted and both the distributor and customer would have been notified, probably through the respective workflow management systems.
21. The customer dispatches a person to the local branch to pick up the left flapper fud. [G-D1.10 and C-S1.2] It, of course, is waiting in the will-call bin when the person arrives at the branch. Upon pick-up, the distributor's order management system completes that line item of the purchase order. [G-D1.11]
22. The left flapper fud arrives at the customer's repair site and is installed together with the widget. [CM6] The maintenance work order is completed and closed out in the CMMS. [CM7] The legacy system interface agent detects that the work order has been completed and closed and triggers an alert

to the Domain Manager. This initiates an AIP to be executed which updates the ontology accordingly and eliminates all the detailed status data that will no longer be needed for analysis purposes.

23. Another AIP will be initiated to transfer the maintenance results information to the distributor so that the data can be included in the Equipment Maintenance Knowledge Base. [KB3] By compiling the completed maintenance information which records what parts were actually replaced and what the general condition of the equipment was both before and after the maintenance, the information will be used in the future to help determine probability of need for future ADNs.
24. That night, the distribution center picks, packs, and ships the widget. [G-D1.9 and G-D1.10]The distributor's logistics systems create the shipping manifest and get the track/trace identification for the shipment. [T-D1.10] This is then passed to the distributor's ontology. The AIP involved requires that the shipping information be sent to the Transportation Agent Server so that shipping status information can be monitored. As shipping status is captured from the carrier, it is passed to the customer since the customer requested a subscription to the status data. The shipping status data is recorded in the ontology. This enables the customer's staff to view the status via the User Interface Web Server if desired.
25. Two days later the widget is delivered to the customer. [C-S1.2] The carrier captures the delivery signature confirmation and records it in its system. [T-D1.11] The Transportation Agent Server will capture the delivery status and form a Tell performative to advise the distributor of the receipt. The status is passed to the distributor's order management system which, in turn, completes the entire purchase order and invoices the customer. [G-D1.13] The customer profile indicates that this customer should be invoiced upon confirmation of delivery.

Admittedly, we started taking great shortcuts in documenting the details of the flow through the various Agent Architecture components in the last part of this example. However, the details covered in the early part of the example can be extrapolated throughout the entire scenario life cycle to provide an overall idea of how this architecture operates for the MRO ISCM.

Page 32 has a diagram of the database design for the equipment knowledge base. This is a very important component of the overall MRO ISCM system. This knowledge base will be operated centrally by Grainger and will provide the data that assists in determining the probabilities that a particular MRO product will be needed when a customer performs a maintenance task on a piece of equipment or components of its facilities.

This corresponds to the architectural element "KB-3 maintenance knowledge Base in the overall SCOR model rendering on page 35 of this notebook.

The knowledge base contains data on generic equipment and facilities as well as on specific equipment. The intent is to have basic maintenance task information available on a generic basis in case the customer does not provide data on a specific piece of equipment. The generic maintenance tasks will include steps plus to the extent possible in a generic situation) a bill of materials for the tools, supplies, consumable and repair parts needed for the task.

Specific instances of equipment will be associated with a generic type of equipment (or facility). The maintenance tasks for the specific equipment will reference their counterparts in the generic area and will specify replacement, deletion, insertion of the steps and MRO parts, bills/materials. To create a maintenance task for a specific piece of equipment, the system will start with the associated generic data (if it exists) and replace, delete and insert details for the specific equipment as indicated.

The bills of material will cite Grainger SKU's whenever possible. They can also cite the applicable product classification parameters and attributes whenever the Grainger SKU does not exist. The class parameter table, class attribute table and attribute value table will be taken from Grainger's product information systems so that these are compatible with the product search engineer used by the Grainger business units.

the intent is to make this data available to customers if they have no order service of maintenance data. Also, the contrasts with the customers should contain provisions for the customers to contribute preventive maintenance data to this repository. Granger will also work with equipment manufacturers and other third party services to make this as comprehensive a knowledge repository as possible.

Actual maintenance history data which includes the use of MRO products will be kept on a specific equipment basis. Where a customer enrolls in the MRO ISCM program, it will register its equipment and facilities inventory with Grainger. The data will be contained in this repository. The knowledge base will also include roll-up records to consolidate dates from multiple instances of the same specific equipment across customers and customer sides. These roll-up consolidation records will have a site code of zero. Equipment and facilities can be comprised of hierarchically organized components. The segment type child table (#99) and the asset child table (#99) provide the definition of these component hierarchies. Maintenance history will generally be compiled at the lowest level of hierarchical detail as appropriate.

The Machinery Information Management Open System Alliance (MIMOSA) is a not-for-profit Trade Association that facilitates the development, dissemination and justification of open, electronic exchange of data for equipment characterization operative, and maintenance information. This organization specifies both data base design schema and data exchange transaction definitions. While MIMOSA is not an industry universal standard, the specifications do embrace a large amount of key features needed in the MROISCM Equipment Knowledge Base. For this reason, I have embraced as much as practical the MIMOSA Specifications in the Equipment Knowledge Base. Only MIMOSA version *1.1* is available as of this date. Version 2 will include more specification that will be useful to this project.

The intent will be to have Grainger agents interact with the customers CMMS/EAM system using MIMOSA transactions if possible and pass data to and from Grainger following (as much as practical) the MIMOSA standards. Where the customers' systems are not MIMOSA-compliant, the Grainger intelligence software agents will translate back and forth between MIMOSA standards and the customer's systems' native interfaces.

This section describes the processing done within Grainger's systems to record new customer equipment in the Equipment Knowledge Base. The process flow charts are attached to page 5.

New equipment is added to the knowledge base in two situations—the first is where a customer is first implemented/installed in the MRO ISCM business program. In this situation, all of the equipment that a customer chooses to embrace within the program is added. The second situation is when an already registered/implemented customer adds a new price of equipment to the program. It should be noted that customers will have a choice of what equipment and facilitator they wish to include in the MROISCM Program. If a customer chooses not to include a specific piece of equipment in the program, that customer can still order MRO products for that equipment/facility but will have to always provide a probability of need in the Advance Demand Notice. If the customer does not include the equipment in Grainger's knowledge base, there will be no way for Grainger to assist in determining the probability of need. Marketing and business policies, rules and procedures will need to be developed for handling MROISCM transactions where the customer's equipment is not included in Grainger's Knowledge Base. This may be a very prevalent case with facilities since many customers do not vigorously catalog their facilities in their CMMS system but rather manage them strictly on a work order basis.

A minimum pre-requisite for a customer to participate in the MROISCM Program is that the customer must use a CMMS (EAM) system to manage the maintenance

TITLE NEW CUSTOMER EQUIPMENT EVENT 18 Project No. ______ Book No. ______

Bob Ramsden 3/29

Robert N Aubert 3-27-00

Jeff Koller 3/30/00

CONFIDENTIAL

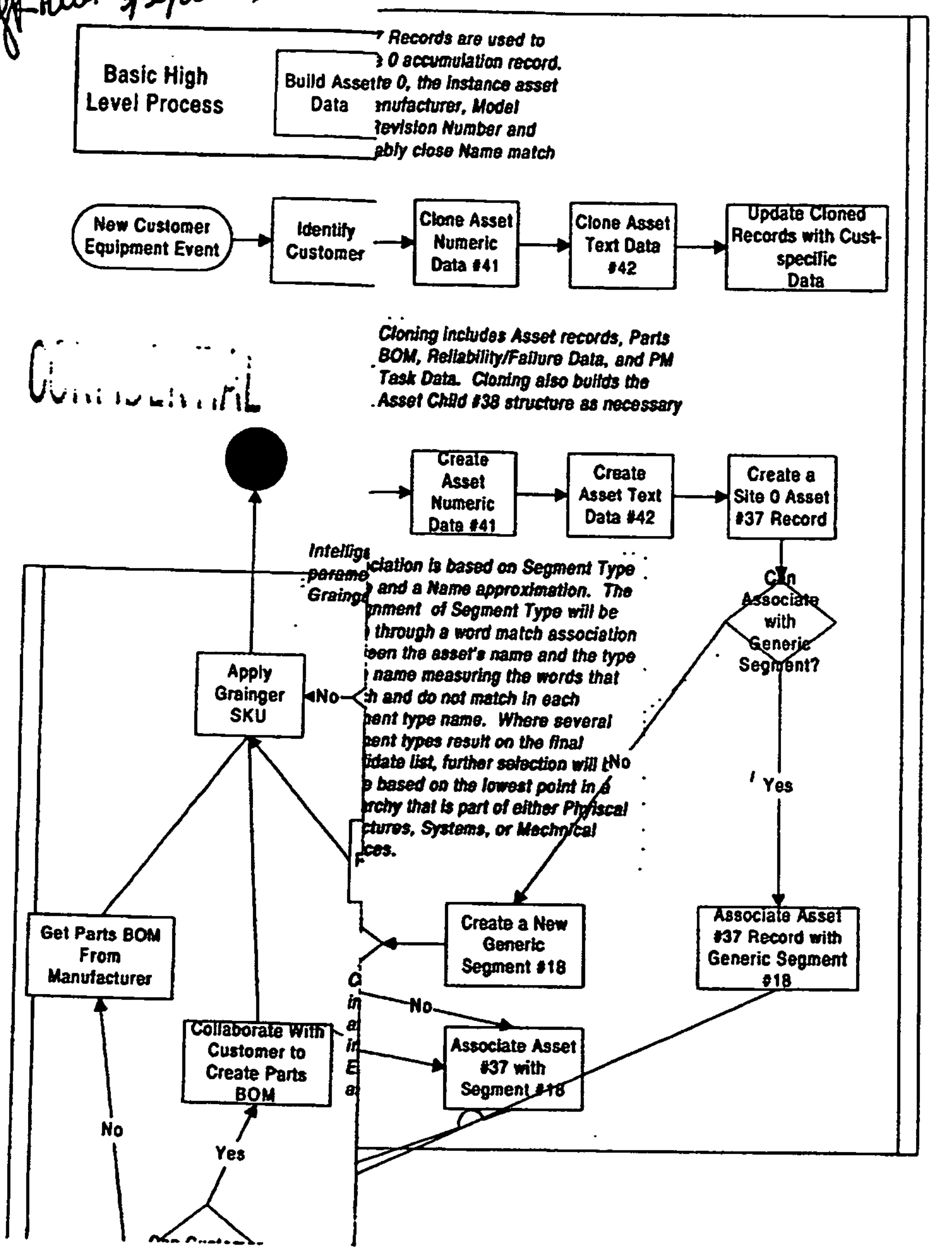

(Page 5 of 35)

Witnessed & Understood by me, Bob Ramsden, Jeff Koller | Date 3/29/00 | Invented by Robert N Aubert | Date 3-27-00

Recorded by

activities. The March, 2000 edition of Maintenance Technology Magazine contains a directory of 40 CMMS/EAM software packages. The major ERP vendors (SAP, Oracle, JDE, etc.) also provide PM functionality although these might not be as robust (best of breed) as the CMMS/EAM packages. Grainger will provide intelligent software agents to interface with these maintenance systems. However, even if a customer uses a CMMS system, the customer does not have to catalogue all equipment and facilities in it.

Some customers may also use equipment condition data acquisition software (example is Canary Labs) to capture and record in a database for analysis equipment condition data from DCS, PLC, and MHI data capture services. These data acquisition systems often supplement the CMMS systems to monitor the condition of the equipment and automatically schedule maintenance tasks when problems are detected. Equipment will be cataloged in these data acquisition systems even if the customer does not catalog the equipment in the CMMS. Grainger will also provide intelligent software agents to interface with these data acquisition systems. When a customer registers in the MRO-ISCM program, parameters are set to indicate which software components the customer is using.

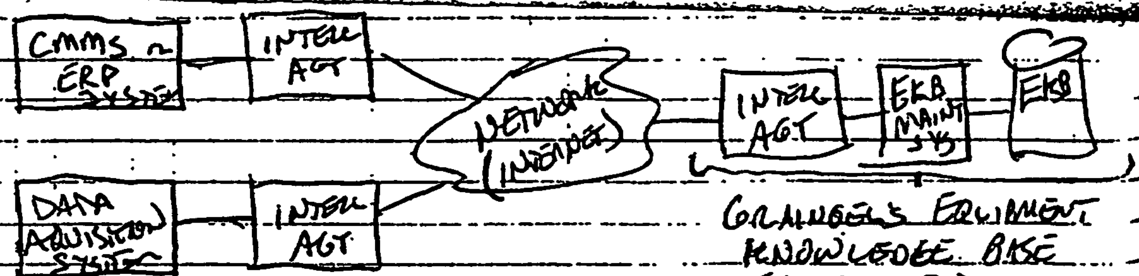

GRAINGER'S EQUIPMENT KNOWLEDGE BASE ENVIRONMENT

To Page No. 21

Witnessed & Understood by me,	Date	Invented by	Date
[signature] [signature]	3/29/00 3/30/00	[signature]	3-22-2000
		Recorded by	

TITLE NEW CUSTOMER EQUIPMENT EVENT Project No.______ Book No.______

From Page No. 20

Robert N Aillen
3-27-2000

Bob Ramsden 3/29/0

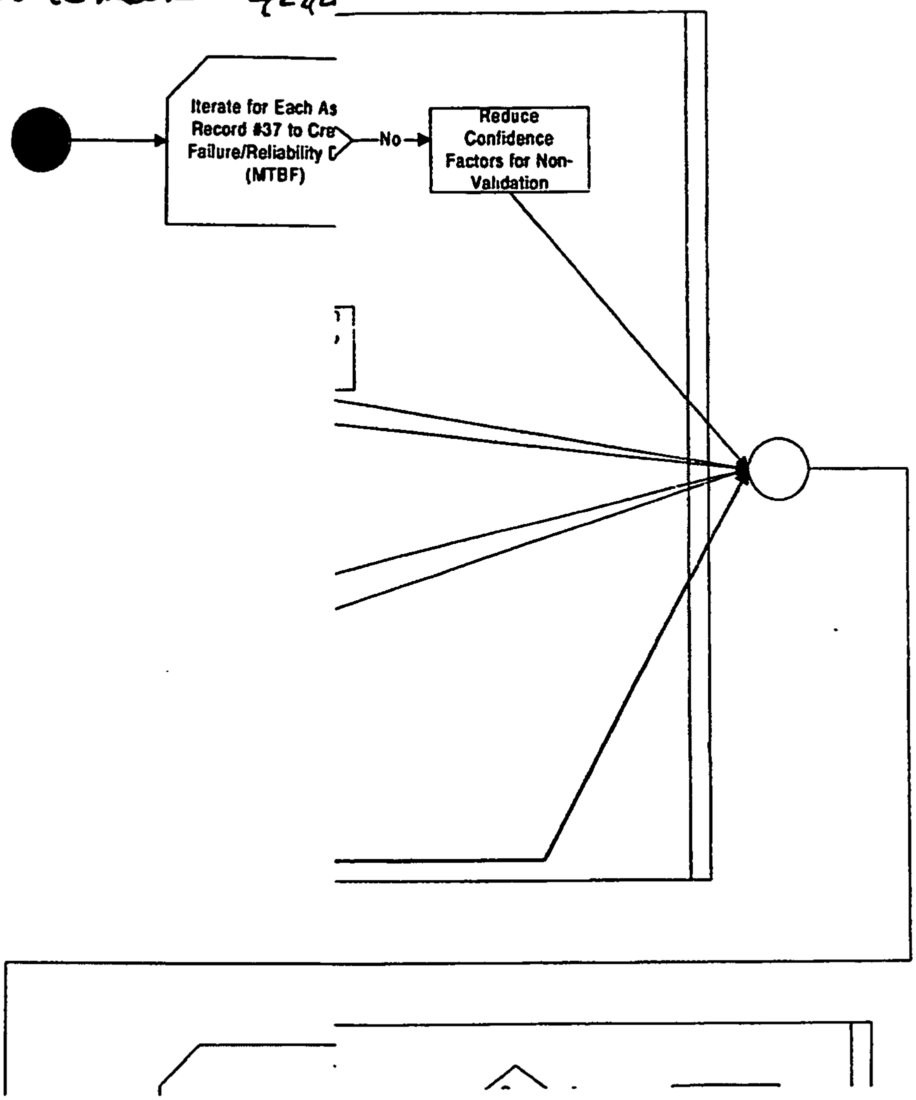

(Page 8 of 35)

22

Witnessed & Understood by me, Bob Ramsden Date 3/29/00

Invented by Robert N Aillen Date 3-27-0 Recorded by

Implied in these process flow charts is that much of the data acquisition from the customers is accomplished by the return of intelligent software agents collaborating with each other to extract the data from the customer's CMMS and/or Data Acquisition Systems. However, involvement will only be needed where the software agents are unable to complete their tasks automatically.

When a customer is first installed and registered in the MRO-ISCM Program, entries about the customer are made in the Grainger Registry (not shown in these process flow diagrams). This will include a list of customer sites with cross reference tables including how a customer identifies a site with how the Grainger Equipment Knowledge Base identifies a site (using a DUNS number ID). Thereafter, the intelligent software agents will be able to automatically translate between the different identification schemes.

When a customer first registers, the software agent on the Grainger side will request the agents attached to the customer's systems to search their data services for information on equipment. For each piece of equipment found by the agents on the customer side, the agents will create a "New Customer Equipment Event" and thus begin the processing depicted in these flow charts.

After the initial registration equipment data load, the software agents at the customer's side will continue to monitor the CMMS or Data Acquisition System to which they are attached for new additions to equipment and factories catalogued in those systems. Where the software agent detects such a new additive, it will work a "New Customer Equipment Event" and pass it to the agent on the Grainger side to initiate the processing flow in these diagrams!

Not shown in these flow charts is a method for the customer to decide not to catalog a piece of equipment or facility in the Grainger Equipment Knowledge Base. Such a designation

will probably be done within the domain of the intelligent software agent(s) attached to the customer's systems. The exact specification of how this will be accomplished will be determined at a later date. The following table shows the implication for a customer not cataloging equipment with Grainger.

CUSTOMER CATALOGS IN CMMS	N	N	N	N	Y	Y	Y	Y
CUSTOMER CATALOGS IN DATA AQUISITION SYSTEM	N	N	Y	Y	N	N	Y	Y
CUSTOMER CATALOGS IN GRAINGER EKB	N	Y	N	Y	N	Y	N	Y
CUSTOMER CAN DO PM PLANNING AND SEND ADN TO GRAINGER	Y	Y	Y	Y	Y	Y	Y	Y
CUSTOMER SUPPLIES PROBABILITY OF NEED DATA	Y	OPTIONAL	Y	OPTIONAL	Y	OPTIONAL	Y	OPTIONAL
GRAINGER PROVIDES SUPPLEMENTARY ANALYSIS OF PROBABILY OF NEED	N	Y	N	Y	N	Y	N	Y
GRAINGER DETERMINES A MORE ACCURATE PROB OF NEED BY COMPARING CUST'S ESTIMATE AGAINST DATA IN EKB	N	N	N	Y	N	Y	N	Y
GRAINGER CAN EXTEND "MOST FAVORED NATION" MARKETING, PRICE, AND PROCESS POLICIES & RULES FOR ADN FULFILLMENT	N	N	N	Y	N	Y	N	Y

To Page No. 24

Witnessed & Understood by me, Bob Romeo [signature] Kelly J Kolle [signature] Date 3/29/00 3/30/00

Invented by Robert N Allen [signature] Date 3-22-00

Recorded by

When the Grainger agent receives a New Customer Equipment Event, the customer system agent may or may not have included a DUNS number for the site ID> However, the customer side agent will have included the site identification taken from the CMMS or Data Acquisition System and this will be correlated to the site—name field in the #13-site record of the EKB. The Grainger agent will send a message back to the customer agent directing the custom-side agent to use the DUNS number for the site.

Depending on how the customer has identified the equipment in the CMMS, the software agents in both the Grainger and customer sides may have to collaborate to resolve any *semantic* difference in manufacturer's name and the equipment names and model number. This will be necessary to determine if a site 0 record already exists for the equipment being added. Once the manufacturer's ID (DUNS code) and the Model Number have been established, then how the customer names the equipment does not need to match exactly the name on the site 0 record. The name on the customer's site detail record for #37 asset will always be the name that is used in the customer's systems. Otherwise, the user tag ID will always be the way the customer's CMMS identifies the specific piece of equipment. Once all of this has initially been established, the customer side software agents will only need to evaluate the DUNS for the site together fee site together with the User-Tag-ID to correctly identify the equipment in the EKB for all subsequent transaction activity (including the ADN).

Where the flow chart steps indicate collaboration is needed to validate something with the customer, the software agents will perform as multi-validation as is possible and then involve all means at the customer site to complete the validation. Validation confirmation and/or changes from the humans will be captured by the customer site agents and sent to Grainger.

Part of the validation of the equipment BOM for parts is to ensure that the BOM in the ERB agrees with the BOM in the customer's CMMS> This may involve some complex collaboration between the software agents and possibly involve the humans for final resolution.

Where validation collaboration with manufacturers is involved, the Grainger systems and network will facilitate the collaboration process, but the actual collaboration will be between humans from the customer and from the manufacturer.

Failure/reliability data in the Grainger EKB will only be at a summary level—detailed only to the degree needed to predict probability of need. This is likely to be some kind of metric time (n order specific) between failures. Detailed observation recording will not be kept by Grainger. This level of detail will only be kept by the customer (either in the customer's CMMS or in the Data Acquisition System). Some customers may choose not to see any reliability data in their systems. If this is the case, Grainger and the customer will have to collaborate to attempt to learn what MTBF metric to use and what the metric value was the last time a particular MRO part was used for the equipment. This will be the minimum data needed for the MRO ISCM to attempt a prediction of probability of need. In the absence of such data, the EKB will capture the first MTBF metric point when the first Maintenance Task is performed and work forward from that point. Also in this case, the confidence factors will need to be set to zero until some basic reference point in the MTBF metric is established.

Once the MTBF data is established for the new equipment, the validation step with the customer will include passing the data on reliability to the customer's agent if the customer's systems will also retain the data. The software agents at the customer side will take care of translation between the reliability data forms and metrics used by the Grainger EKB and the forms and metrics used by the customer's systems.

An equipment retirement is triggered when a customer removes a piece of equipment from its CMMS or Data Acquisition System. The Grainger intelligent software agent that interfaces with the customer's CMMS or DAS will detect this deletion action and indicate the transaction to the Grainger system.

If the customer does not catalog the equipment in a CMMS or DAS system, then the customer will need to notify Grainger of the equipment retirement by contacting Grainger's customer service. The customer service group will then initiate the system transaction for this event.

Equipment asset records (#37 and associated subordinate records) will not actually be detected from the EKB from this event. Instead, the record's status flag will be set to detection: The MIMOSA standards specification has a data element called "RSTAT_TYPE_CODE" in every table. The suggested values for this new status code arc:

1 = Active row

2 = Inactive Row

3 = Soft deleted row

The Customer Retire Equipment Event will cause the row status code to be set to "2" to indicate the row is inactive. This enables all the historical data (especially the reliability data) to be retained and associated with the Site 0 roll-up record for all instances of that particular equipment. The intelligent agents which are predicting probability of need for an ADN for other instances of the same equipment will likely need to refer to the historical data about the retired piece of equipment.

Also, the piece of equipment may subsequently be installed at a different site of the customer's or be sold to a different customer. In this case, where the New Customer Equipment Event is processed, the maintenance programs will find the evidence record by matching in manufacturer's DUNS code, Model Number, Revision Number, and Serial Number. The new Asset Off-site number, User Tag ID and (possibly) name fields can be updated in the #37 Asset record, the row status re-set to 1 = active row, and all of the subordinate records (maintenance history, parts BOU, MTFB data will be carried forward. The actual PM tasks may have to be revised and validated by the new owner/site.

If a site O roll-up asset record has all its instance records' status set to inactive, then this indicates a situation where there are no longer any instances of Grainger customers in the MRO-ISCM program who are using their equipment. When this happens, the row status should be set to 3 = soft delete. Then the entire set of affiliated records will be subject to some kind of records retention policy. This records retention policy will indicate the period of time before the records will be physically deleted from the EKB data base management systems.

When the physical deletion actually occurs, this record will be placed in an archives file.

Part of the records retention policy will specify that the equipment manufacturers should confirm that the equipment is truly outdated and no longer being used by any of its customers.

The process flow chart on page 17 depicts the overall steps needed to record maintenance results in the Equipment Knowledge Base (EKB).

Capturing and recording maintenance history data is truly a collaborative effort between the customers' and Grainger's systems. The maintenance data can only originate at the customer's site. The purpose of capturing maintenance data is to create in the Equipment Knowledge Base data about the reliability of the MRO products used to maintain the customer's

equipment. This reliability data is used to predict the need for a specific MRO product, where some maintenance job is performed on the equipment. By predicting the probability of need, safety stocks throughout the entire MRO-ISCM supply chain can be better managed and minimized. The minimization of safety stock throughout the entire supply chain is a major business objective of MRO-ISCM.

There are two ways that maintenance data can be captured and sent to Grainger for inclusion in the EKB. The first is when a customer places PO for the MRO product. If the PO can be associated with the maintenance tasks, then the items in the PO can be analyzed for inclusion in the EKG. Note that not all of the MRO product used in the maintenance task will show up on a PO sent to Grainger. If the repair part was obtained from the customer's own safety stock (MRO inventory storeroom or tool crib), then no PO to Grainger was issued. At some later time, the customer may issue a PO to Grainger to replenish its DON MRO inventory stock, but it is not likely that that replenishment PO will always be traceable to some maintenance action.

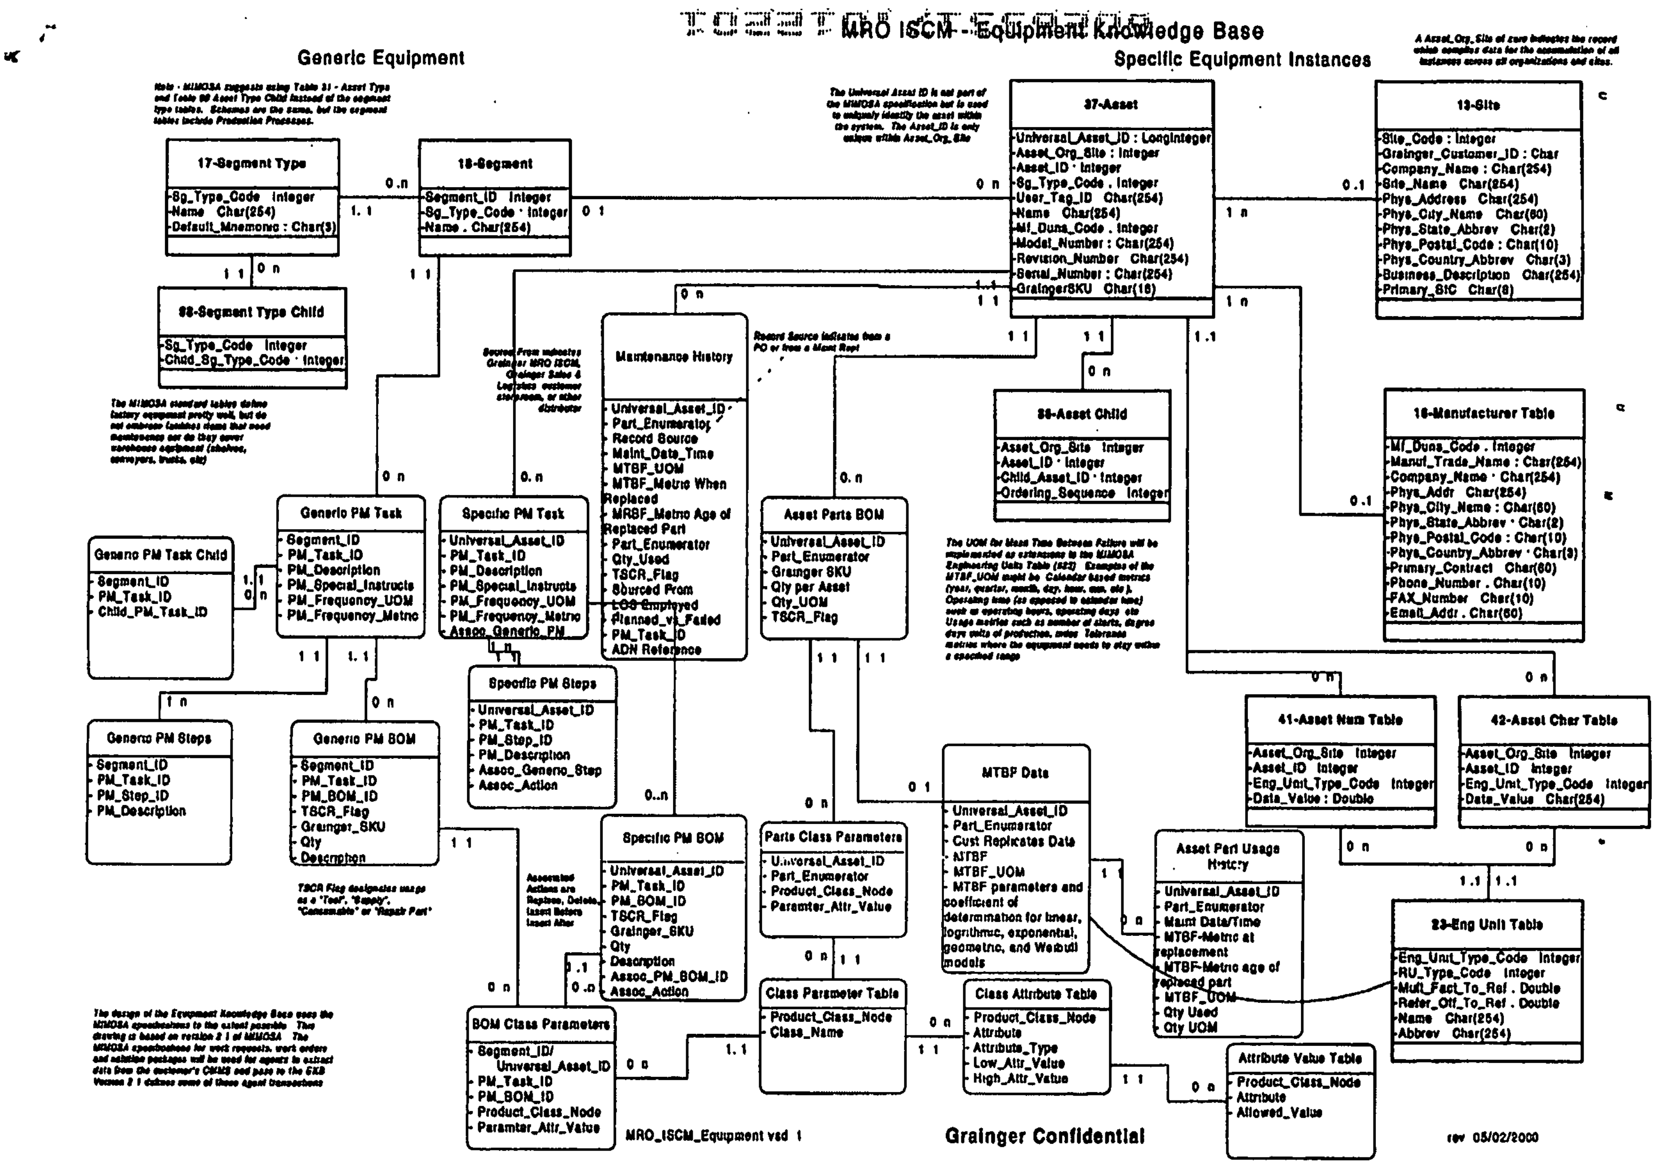
MRO ISCM - Equipment Knowledge Base
Generic Equipment
Specific Equipment Instances
17-Segment Type
Sg_Type_Code Integer
Name Char(254)
Default_Mnemonic : Char(3)
18-Segment
Segment_ID Integer
Sg_Type_Code · Integer
Name . Char(254)
88-Segment Type Child
Sg_Type_Code Integer
Child_Sg_Type_Code · Integer
37-Asset
Universal_Asset_ID : LongInteger
Asset_Org_Site : Integer
Asset_ID · Integer
Sg_Type_Code . Integer
User_Tag_ID Char(254)
Name Char(254)
Mf_Duns_Code . Integer
Model_Number : Char(254)
Revision_Number Char(254)
Serial_Number : Char(254)
GraingerSKU Char(16)
13-Site
Site_Code : Integer
Grainger_Customer_ID : Char
Company_Name : Char(254)
Site_Name Char(254)
Phys_Address Char(254)
Phys_City_Name Char(60)
Phys_State_Abbrev Char(2)
Phys_Postal_Code : Char(10)
Phys_Country_Abbrev Char(3)
Business_Description Char(254)
Primary_SIC Char(8)
16-Manufacturer Table
Mf_Duns_Code . Integer
Manuf_Trade_Name : Char(254)
Company_Name · Char(254)
Phys_Addr Char(254)
Phys_City_Name : Char(60)
Phys_State_Abbrev · Char(2)
Phys_Postal_Code : Char(10)
Phys_Country_Abbrev · Char(3)
Primary_Contract Char(60)
Phone_Number . Char(10)
FAX_Number Char(10)
Email_Addr . Char(60)
38-Asset Child
Asset_Org_Site Integer
Asset_ID · Integer
Child_Asset_ID · Integer
Ordering_Sequence Integer
Maintenance History
Universal_Asset_ID
Part_Enumerator
Record Source
Maint_Date_Time
MTBF_UOM
MTBF_Metric When Replaced
MRBF_Metric Age of Replaced Part
Part_Enumerator
Qty_Used
TSCR_Flag
Sourced From
LCG Employed
Planned_vs_Failed
PM_Task_ID
ADN Reference
Generic PM Task
Segment_ID
PM_Task_ID
PM_Description
PM_Special_Instructs
PM_Frequency_UOM
PM_Frequency_Metric
Specific PM Task
Universal_Asset_ID
PM_Task_ID
PM_Description
PM_Special_Instructs
PM_Frequency_UOM
PM_Frequency_Metric
Assoc_Generic_PM
Asset Parts BOM
Universal_Asset_ID
Part_Enumerator
Grainger SKU
Qty per Asset
Qty_UOM
TSCR_Flag
Generic PM Task Child
Segment_ID
PM_Task_ID
Child_PM_Task_ID
Generic PM Steps
Segment_ID
PM_Task_ID
PM_Step_ID
PM_Description
Generic PM BOM
Segment_ID
PM_Task_ID
PM_BOM_ID
TSCR_Flag
Grainger_SKU
Qty
Description
Specific PM Steps
Universal_Asset_ID
PM_Task_ID
PM_Step_ID
PM_Description
Assoc_Generic_Step
Assoc_Action
Specific PM BOM
Universal_Asset_ID
PM_Task_ID
PM_BOM_ID
TSCR_Flag
Grainger_SKU
Qty
Description
Assoc_PM_BOM_ID
Assoc_Action
Parts Class Parameters
Universal_Asset_ID
Part_Enumerator
Product_Class_Node
Paramter_Attr_Value
MTBF Data
Universal_Asset_ID
Part_Enumerator
Cust Replicates Data
MTBF
MTBF_UOM
MTBF parameters and coefficient of determination for linear, logrithmic, exponential, geometric, and Weibull models
Asset Part Usage History
Universal_Asset_ID
Part_Enumerator
Maint Date/Time
MTBF-Metric at replacement
MTBF-Metric age of replaced part
MTBF_UOM
Qty Used
Qty UOM
41-Asset Num Table
Asset_Org_Site Integer
Asset_ID Integer
Eng_Unit_Type_Code Integer
Data_Value : Double
42-Asset Char Table
Asset_Org_Site Integer
Asset_ID Integer
Eng_Unit_Type_Code Integer
Data_Value Char(254)
23-Eng Unit Table
Eng_Unit_Type_Code Integer
RU_Type_Code Integer
Mult_Fact_To_Ref . Double
Refer_Off_To_Ref . Double
Name Char(254)
Abbrev Char(254)
BOM Class Parameters
Segment_ID/ Universal_Asset_ID
PM_Task_ID
PM_BOM_ID
Product_Class_Node
Paramter_Attr_Value
Class Parameter Table
Product_Class_Node
Class_Name
Class Attribute Table
Product_Class_Node
Attribute
Attribute_Type
Low_Attr_Value
High_Attr_Value
Attribute Value Table
Product_Class_Node
Attribute
Allowed_Value
TSCR Flag designates usage as a "Tool", "Supply", "Consumable" or "Repair Part"
Record Source indicates from a PO or from a Maint Rept
MRO_ISCM_Equipment vsd 1
Grainger Confidential
rev 05/02/2000

(Page 17 of 35)

The second way to capture maintenance results data is to have the maintenance staff complete maintenance documentation, either as the maintenance work progresses or immediately after it is completed. In many cases, this takes a good deal of discipline on the part of the customer's management to ensure that the maintenance staff does, in fact, document its work. However, companies are growing in the awareness of the importance of well-managed maintenance programs and having well-documented maintenance history is an acknowledged best practice.

THE MRO-ISCM program will utilize both methods of capturing this data. This is primarily as a check and balance—if the data does not come through on one channel, it hopefully will come through on the other channel. And if it came through on both channels—so much the better!

If the ADN indicated a 100% probability/need, and the company's policy was to order against the ADN rather than issuing a separate PO, then a virtual consumption PO against the ADN will be created within Grainger's systems. If an ordinary PO is issued by the customer, then it must reference some data that can be used to match the PO to the maintenance of specific equipment. That matching logic will be discussed in greater detail in a separate section of the research.

The PO's will be processed through Grainger's order taking systems and eventually pop out to the system functionality that maintains the EKB. The exact architecture of this process and interfacing with EKB systems is beyond the scope of this present research. It may involve Business Event Engineers, interfaces with SAP R/S and/or interfaces with back office legacy systems.

Once the association is made between the PO and maintenance in a specific piece of equipment, the system needs to determine if a maintenance history record has been previously created as a result of a maintenance results document. If one exists, the reconciliation process will be triggered. If one does not exist, a new maintenance record will be added with as much data as is available from the PO. The Record-Service field will be set to indicate this data came from a PO rather than from a maintenance document.

If an associated maintenance document transaction does not come within some period of time, then the system will try to complete the maintenance history record with the available data. If not enough valid data is available, agent collaboration can query the customer to solicit the missing data. The agent at the customer site will likely need to query a human for the data if it cannot extract it out of the customer's CMMS or DAS system. Once enough data has been captured from the customer to make a valid maintenance history record, the system will proceed to re-compute reliability statistics for the equipment.

The maintenance staff can use a variety of methods and tools to create documentation on maintenance actions taken. These can range from totally non-technical paper-based documentation reporting to the other extreme of state-of-the-art multimedia-based PDA/HPC or wearable computers that are RF-Connected to hot systems and allow maintenance staff to voice dictate their maintenance curb as it progresses. The flow diagram on page 17 indicates several of these possibilities. The on-site PC's are typically PC's located on the production floor near the equipment being maintained. Or they may be PC's located in the maintenance office where the maintenance staff can return to file their reports. Maintenance staff may also use laptop PC's which may or may not be connected to the customer's LAN while maintenance is being performed. In this situation, the maintenance person sues the laptop to record maintenance

actions, and through some other means, this data is subsequently transferred from the laptop to the other systems.

Currently very popular is the use of PDA's and HPC's. Here, Windows CE is a dominant operating system, although the PACMOS also enjoys considerable market share. These hand-held devices may be connected to host systems either through a wire-based media or through RF connection, or they may be completely standalone and requiring subsequent sign-up with a PC connected to a LAN> Here again, the maintenance person will use the hand-held device to record the maintenance tasks. See Microsoft Speech API.

Advances in speech recognition technology now make digital dictation, hand-held devices a viable option for this data collection. Both Lemout and Hauspie and Dragon systems offer low-cost (between $200 and $300 single-visit pricing) units that permit dictated maintenance results to be automatically transcribed (in digital form) into host applications. In this case, there will be a transcription agent that will be responsible for re-formatting and inserting the dictated maintenance data into the destination host system. This transcription agent (likely to be provided to the customers by Grainger) will need to embrace nature language processing technologies to symmetrically interpret the maintenance person's dialogue into the structural data needed by the destination host systems. Given the low cost and ease of use, their approach from the perspective of the maintenance person, this may be a highly popular approach. The maintenance person dictates the maintenance documentation into the hand-held device. The device is synced to a PC that has the transcription agent software, and the agent software transfers the structural data to the destination host system.

At the extreme hi-tech end of the spectrum is the multimedia HPC connected via RF to a host CMMS system. Vocollect, Inc. offers a windows CE-based wearable computer that

falls into this category. In this case, the maintenance person is connected in real time to the CMMS (and perhaps order performance support systems to use the multimedia technology to guide the person through complex diagnostics and maintenance tasks). Voice input is the primary media for these devices. The maintenance person speaks answers to audio or visual prompts from the device (driven by the host systems) and can also dictate the results of the maintenance tasks. As with digital dictation, some kind of transcription agent is likely to be needed.

Since most customers do not make heavy technology investments in the maintenance function, we can assure that a large number of customers will use the lower-tech approaches. Grainger will need to offer intelligent software agent technology to support all data collection approaches.

Collecting this data is an extremely important aspect of the MROISCM concept, so extreme care must be directed to designing the systematic and human support for these functions. In order to ensure that the right data gets collected, the maintenance person will need to follow scripts that solicit all the necessary pieces of data. If the data is ultimately going through a CMMS system, then the CMMS system may have such scripts embedded in the design of its user interface screens. In other cases, Grainger query have to provide these scripts. These may be in the form of low-tech "cheat sheets" used for written documentation, pre-printed written form data collection programs for the PC's, PDA's/HRC's etc. and guideline prompts for digital dictation.

At a minimum, Grainger will need the following data to populate the EKB:

- Identification of the customer, side, and equipment
- Date the maintenance's performed

- The current usage, meter reading, or some other meter used to track reliability and failure – also life of replaced part.
- The MRO product(s) used and the associated quantities
- Whether used as tools, supplies, consumable, or a repair part, it would also be used to capture the service of the MRO products – Customer's MRO inventory stores,

Insert # 3

RECORD MAINTENANCE RESULTS

IN THE EQUIPMENT KNOWLEDGE BASE

the Grainger MRO-ISCM program, other Grainger procurement channels. (save as Branebas or Web site), other MRO distributors.

Regardless of how the data is collected, it will end up in one of three destination systems.

- Many CMMS systems are designed to capture the maintenance results and history data
- Some customers may use a Data Acquisition System to track-the-condition/equipment, and some of these systems also are designed to collect maintenance data
- Some customers may use neither. In this case, Grainger will provide a maintenance action capture agent software system.

Grainger will provide interfacing agents software for the CMMS/EAM and the DAS systems. These agents will provide for the addition of new maintenance results data, extract the data reformat and translate as needed and pass the date to Grainger.

When the data is received by Grainger, the ERB maintenance system must first determine if the same event data has already been received the way the purchase order channel. If it has, then a reconciliation process is started. If it has not, then a new Maintenance History Record is created and added to the data base.

The reconciliation process will need to be intelligent to reconcile things like PO dates being different from maintenance dates (the maintenance date is the fixed data used in the Maintenance history record.) Looking for duplicate PO's or DO qualifier which included customer storeroom replenishment as well as MRO repair path used in the maintenance, and so forth.

It is possible that the MRO product cited in the maintenance history documentation has not yet been included tin the Asset Path BOM. In this case, a new Asset Path BOM record must be added or the specific equipment at this customer site and also for the Site O equipment accumulation record.

Once the maintenance history record has been added to the EKB, the reliability/failure status needs to be recomputed for the individual price of equipment and for the Site O record. Note that reliability states will not be computed for "TOOLS" but will be for "REPAIR PARTS." Usage status for consumables and supplies will become the basis for subsequent prediction of the probability and need when an ADN is processed, and ultimately provide a basis for redeeming safety stub levels throughout the supply chain.

Reliability data in the Equipment Knowledge ASE (ERB) is kept in order to predict probability of need for an MRO product when an Advance Demand Notice (ADN) is received that cites maintenance work to be done on a piece of equipment for which there is data in the ERB.

TOOLS – Reliability data will not be kept for MRO, product classified as tools. Tools are supposed to be used multiple times across many maintenance jobs.

SUPPLIES – Data on supplies will be kept in the form of quantity needed per maintenance task. Supplies are typically not re-used in multiple maintenance tasks but are disposed of when the task is completed. Some times this quantity will be fixed. Other times it may shovel around some kind/Bell Curve. The MTBF data record for the supply asset part will contain the mean and standard derivation will be computed from the Maintenance History records for the asset path usage. (See STATS function in c/MATH). When the ADN cites a supply SKU, the EKB will return the probability for the cumulative distribution in the quantity on the ADN. The probability formula

Q = quantity

$$P(Q) = \left(\frac{1}{(\text{std dev})\sqrt{2\pi}}\right) e^{-\frac{1}{2}\left(\frac{Q-\text{Mean}}{\text{std dev}}\right)^2}$$

alternative is to standardize

$Z = (Q - \text{Mean})/(\text{std dev})$ and then use normal function

To Page No. 39

Witnessed & Understood by me,

[signatures]

Date 4/3/00 4/4/00

Invented by Robert N Dellen

Date 3-31-00

Recorded by

<u>CONSUMABLES</u> – Data on consumables will kept in the form of quantity consumed in a period of time or time-the unit/measure like production units. The Asset Part Usage History Record will show the quantity used and the MTBF – Metric for age of part will be the "time" since the consumable was last replenished in the equipment. This data will also be kept as a normal distribution with a mean and standard derivation. When an ADN is received, it will cite the quantity of consumable expedited to be used and the time since last replenishment. The system will convert to quantity/time and then compute P(N) where X is the quality per time result.

Since both Supplies and Consumables use the normal distribution, when the supply quality or the consumable quantity per time equal their respective means, the probability will be 50%. The system will add 50% so that the MRO ISCM probability will be 100% of the anticipated demand is for the mean quantity (or quantity per time).

In some cases (especially in tools) the standard derivative will be zero (ex: always need 1 hammer). In this case the probability that the ADN quantity will be needed will always be 100%.

In most cases, when customers list tools and/or consumables on the ADN, the customer will be designating 100% probability of need, and what results from the ERB will generally not matter.

<u>REPAIR PARTS</u> – For repair parts, the asset part usage. History will contain records of when specified repair parts are replaced in the equipment. The history record will show the "time" when the part was replaced and the "age" of the replaced part. If the user does not specify the age of the replaced part, the system can figure it out from the previous history records unless this is the first record.

The reliability data will be kept in the form failure "ages". When the ADN is received, it will show what the current "age" of the existing part and the system will compute the probability of failure (need) for the given "age" of the current part, the system can compute it from prior history records in the EKB unless their is no prior history of replacing this part.

The Weibill distribution will be the preferred distribution curve. However, the system will also compute best fit curves for linear curves, logarithmic curves experimental curves and geometric power curves. The curve fitting programs all compute K^2 = coefficient of determination (or correlation coefficient). When the ADN is processed, the system will use the type curve that has the best R^2 (that is the greatest R^2 value – $R^2 = 1$ is best fit and $R^2 = 0$ is essentially no fit).

The Weibill distribution function include the Reliability Function R(t). The probability failure is 1 – R(t) and this is also the probability of need. The history of usage (failure) will contain the failure data for the part – the "age" of the part replaced. Weibill curve fitting will compute the B = shape, 0 = scale or characteristic life, and 5 = location parameters along with the associate R^2 correlation coefficient. If there are less than 7 sample data records, the Anaid Plot algorithm will be used. For samples 7 or greater, the Probability Plotting Algorithm will be used. Some references claim that the Maximum Likelihood Estimation. (MLE) algorithm works best for sample sizes greater then 15. However, MLE does not compute an R^2 value. It does not appear that the results are that different between Probability Plotting (PP) and MLE – especially when B gets larger. Most MRO products will have larger B parameters. Therefore, for consistency in computing and using an R^2 value, PP algorithm will be used for all sample sizes 7 or greater.

For the four "linear" regression models, the data will represent probability of failure at age t

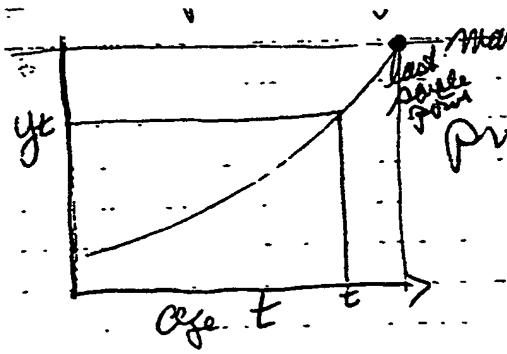

Probability of failure $= \dfrac{y_t}{y_{max}}$

where $y_m = f\left(\begin{array}{c}\text{the oldest } t\\ \text{in file}\end{array}\right)$

$y = f(t)$

To Page No. 42

Witnessed & Understood by me,	Date	Invented by	Date
[signatures]	4/3/00 4/4/00	[signature]	3-31-00
		Recorded by	

Weibill censor data will not be considered in the EKB. MRO ISCM is not really about product failure but about product replacement in the equipment or need in a maintenance job. This concept of sensorial data is not relevant for MRO ISCM. All data will be considered "complete" for purposes of using the Weibill algorithm.

The Asset Part Usage History record replaces the MTBF Probability record in the data structure diagram on page 15. The Asset Part Usage History will have the following data elements:

- Universal Asset ID
- Part Enumerator
- Maintenance Date and Time
- MTBF Metric at time of use
- MTBF Metric for age of repair part replaced or time deemed last consumable was replenished.
- MTBF UOM
- Quantity Used
- Quantity UOM

MIMOSA table # 23 is the Engineering Unit Type table and is the table of unit of measures for the EKB.

For each UOM, there is a conversion to other UOM's within the same category. The category is designated by the data element RO Type. For example, RO Type 8 is weight, RU Type 10 is for distance, etc.

Each UOM record has a factor and offset to convert to a reference UOM for the category:

Therefore, to convert from 4 units (UOM=a) to y units (UOM = B)

•

$$\frac{4\text{-offset (a)}}{\text{factor (a)}} = \text{ref unit} = \frac{\text{y-offset (b)}}{\text{factor (b)}}$$

$$y = \frac{(\text{x-offset (a)}}{\text{factor (a)}} * \text{factor (b)} + \text{offset (b)}$$

Ex: 3 inches = ¼ ft. to convert 4 = 3 inches to y ft: (see next page)

$$y\ (\text{ft}) = \frac{(3\ \text{in} - 0\ (\text{offset (a)})}{(\ 39.4\ (\text{factor (a)})} \underset{\text{factor B}}{*\ 3{,}280{,}834} \underset{\text{offsets B}}{+\ 0}$$

=.24,98 which rounds to 0.25 ft.

In version 1.1 of MIMOSA table # 23 the data for conversion to °F should be – 6453.9232.

In many instances the data in version 1.1 is not given at much precision. We may need to edit the file to add more precision to the conversion data.

Also, the data element for abbreviating the UOM should also be included if it ever becomes necessary to convert from UOM abbreviation to EV Type which is the key field in table # 23.

Additional entries will be needed in this table to accommodate units product, equipment cycles, on of cycles, etc. in order to used the table for UOM's tat measure "age" for MRO parts.

Most of the new CRIS model from MIMOSA involves Work Requests, Work Order and Solution Packages that will be maintained by the customer's CMMS/EAM system. The MIMOSA standards also include minimal specifications for intelligent agents interacting with a MIMOSA compliant database and the associated XML DTD's used in the exchange of data between agents. This will be helpful later when we start looking at the agent interacting with the customer's systems.

New Customer Equipment Data Added to Equipment Knowledge Base

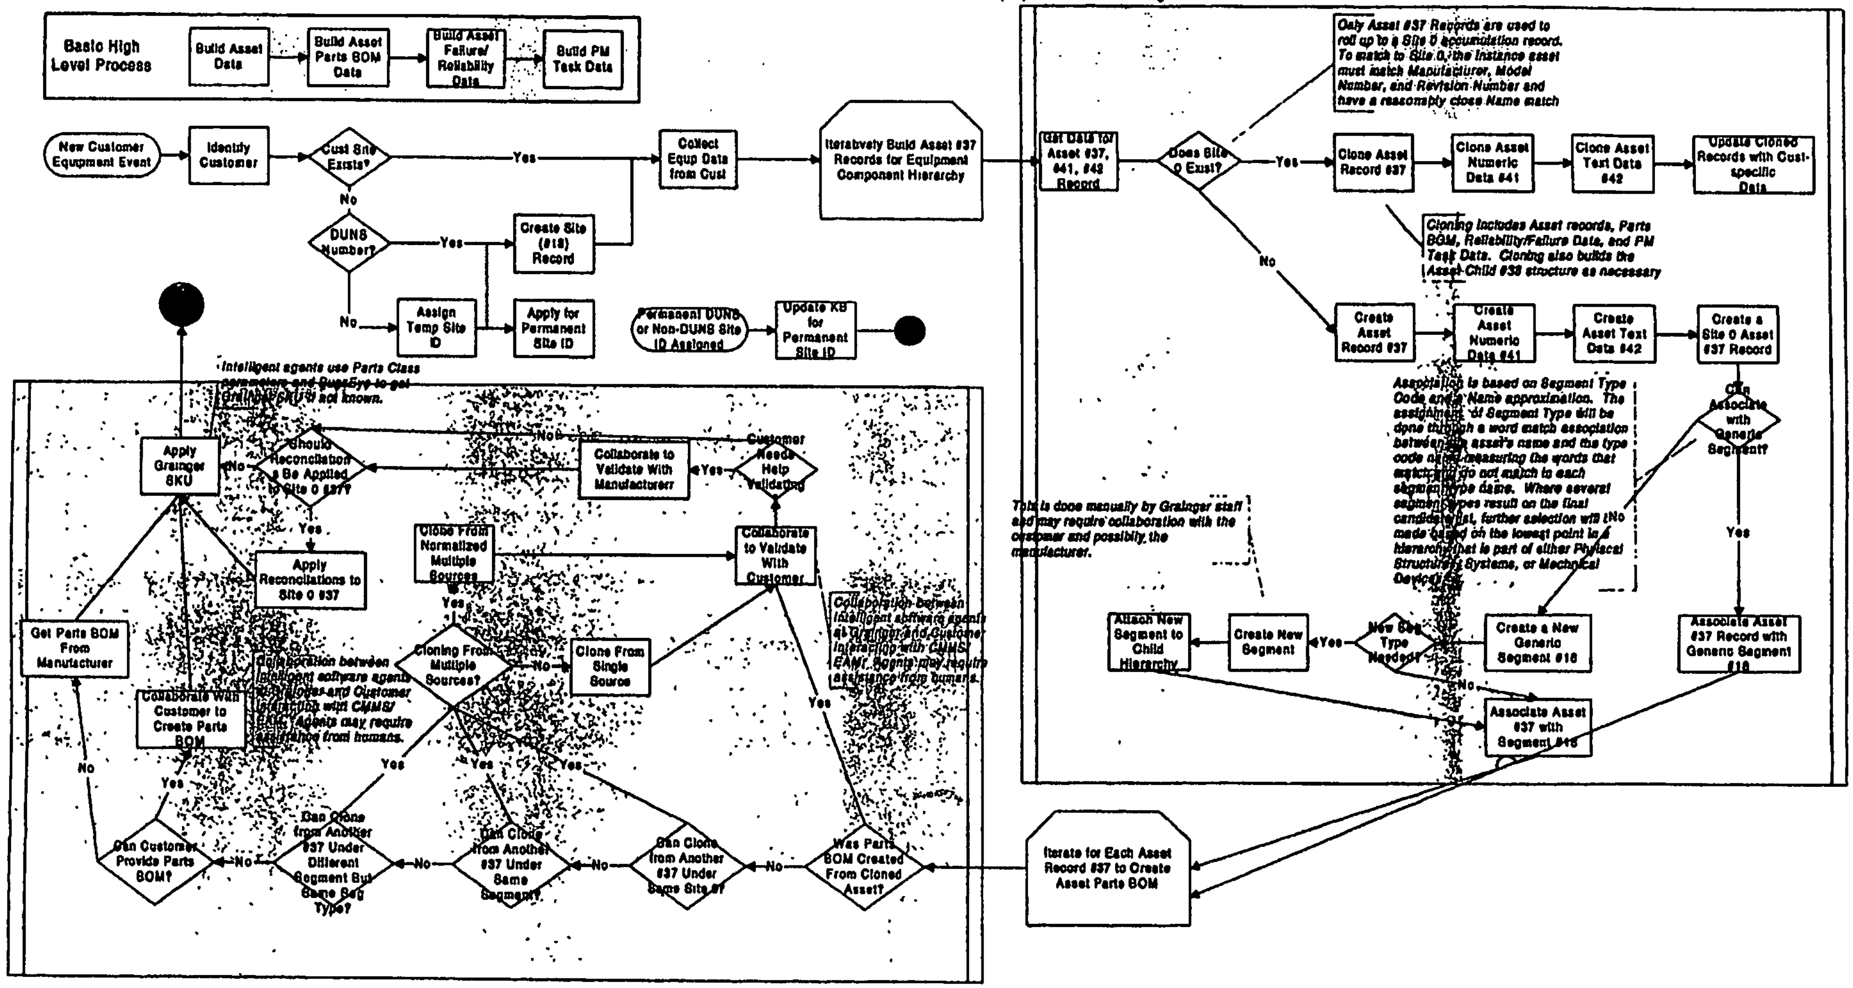

(Page 32 of 35)

New Customer Equipment Data Added to Equipment Knowledge Base (cont)

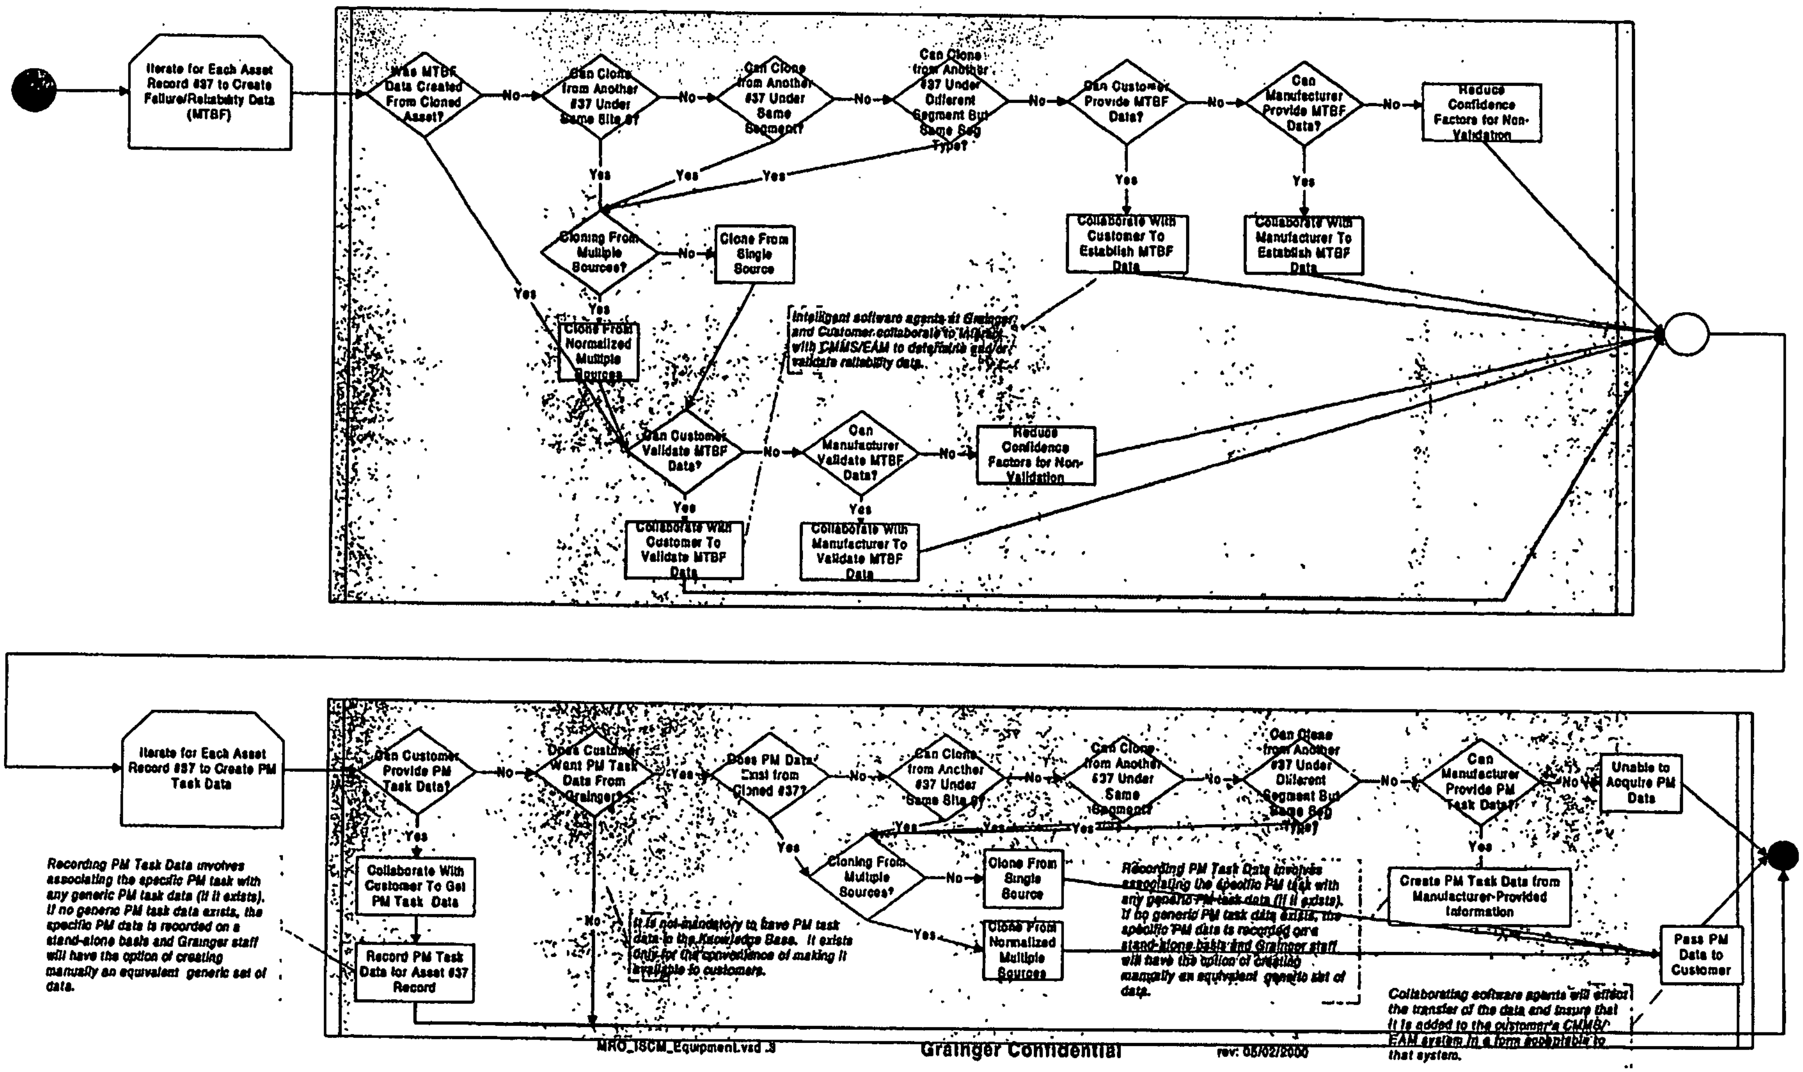

(Page 33 of 35)

Record Maintenance Results in the Equipment Knowledge Base

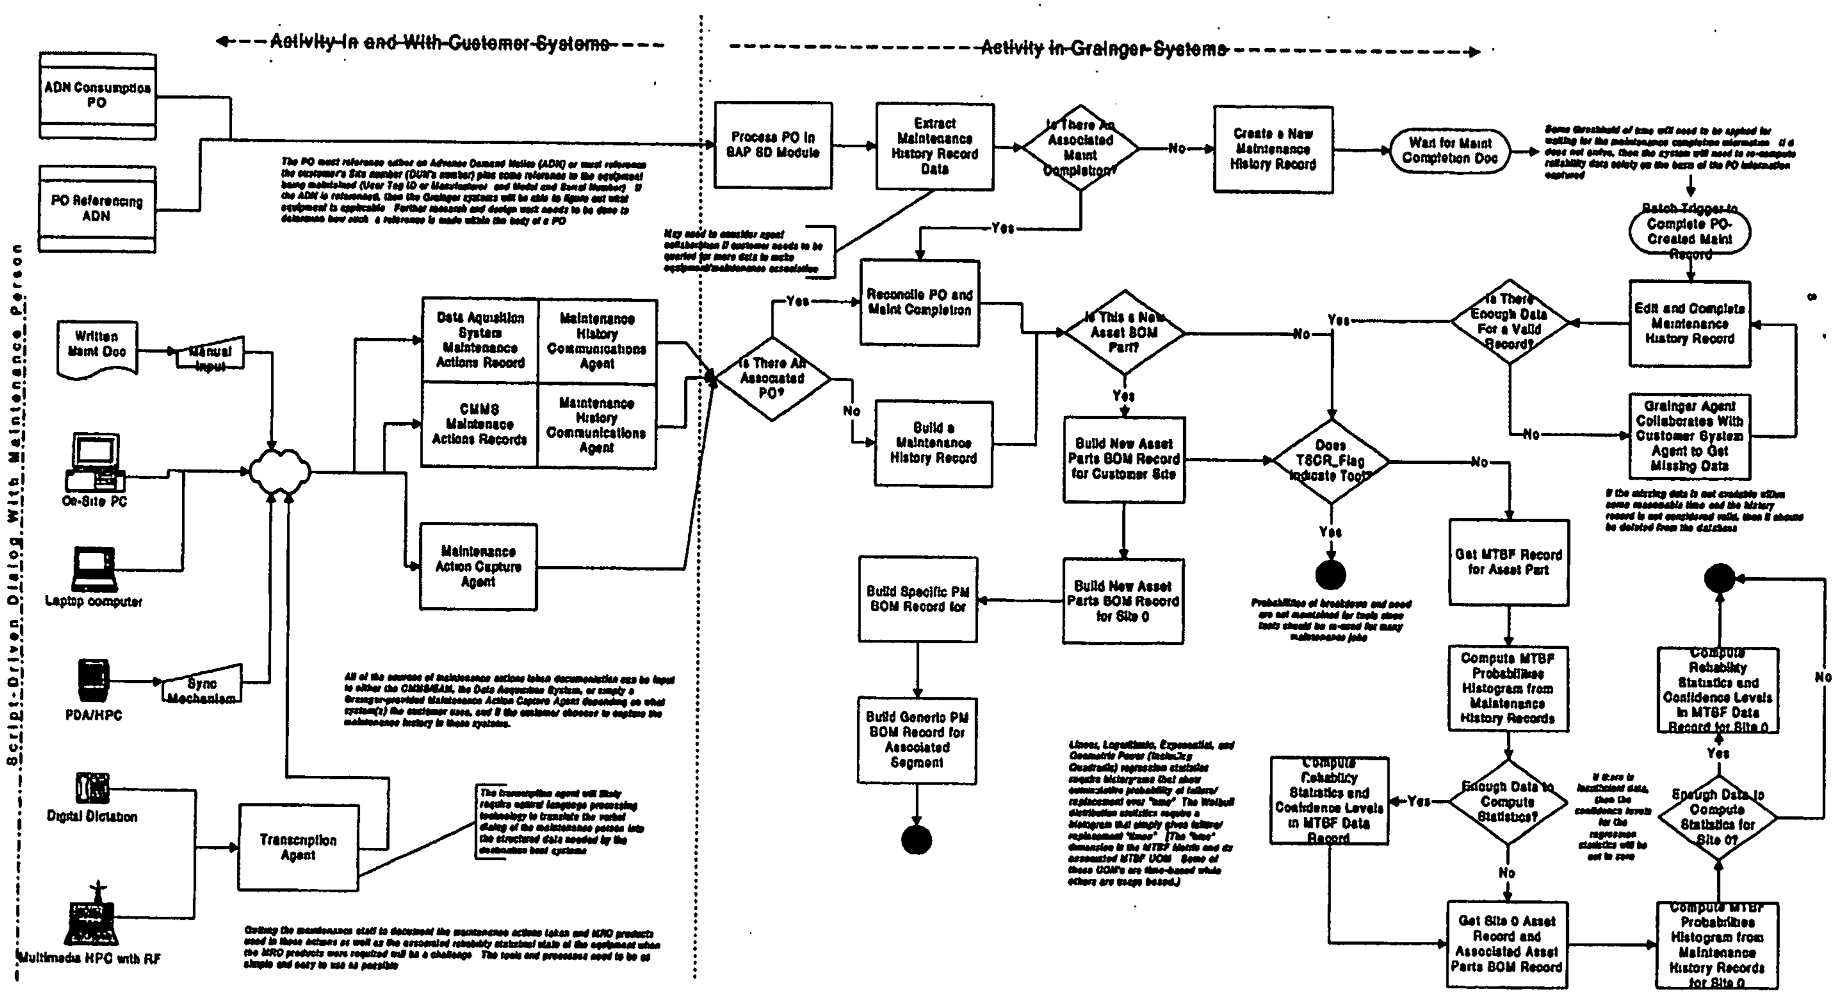

(Page 34 of 35)

TITLE

Project No.______

Book No.

SCOR Model Rendering

CONFIDENTIAL

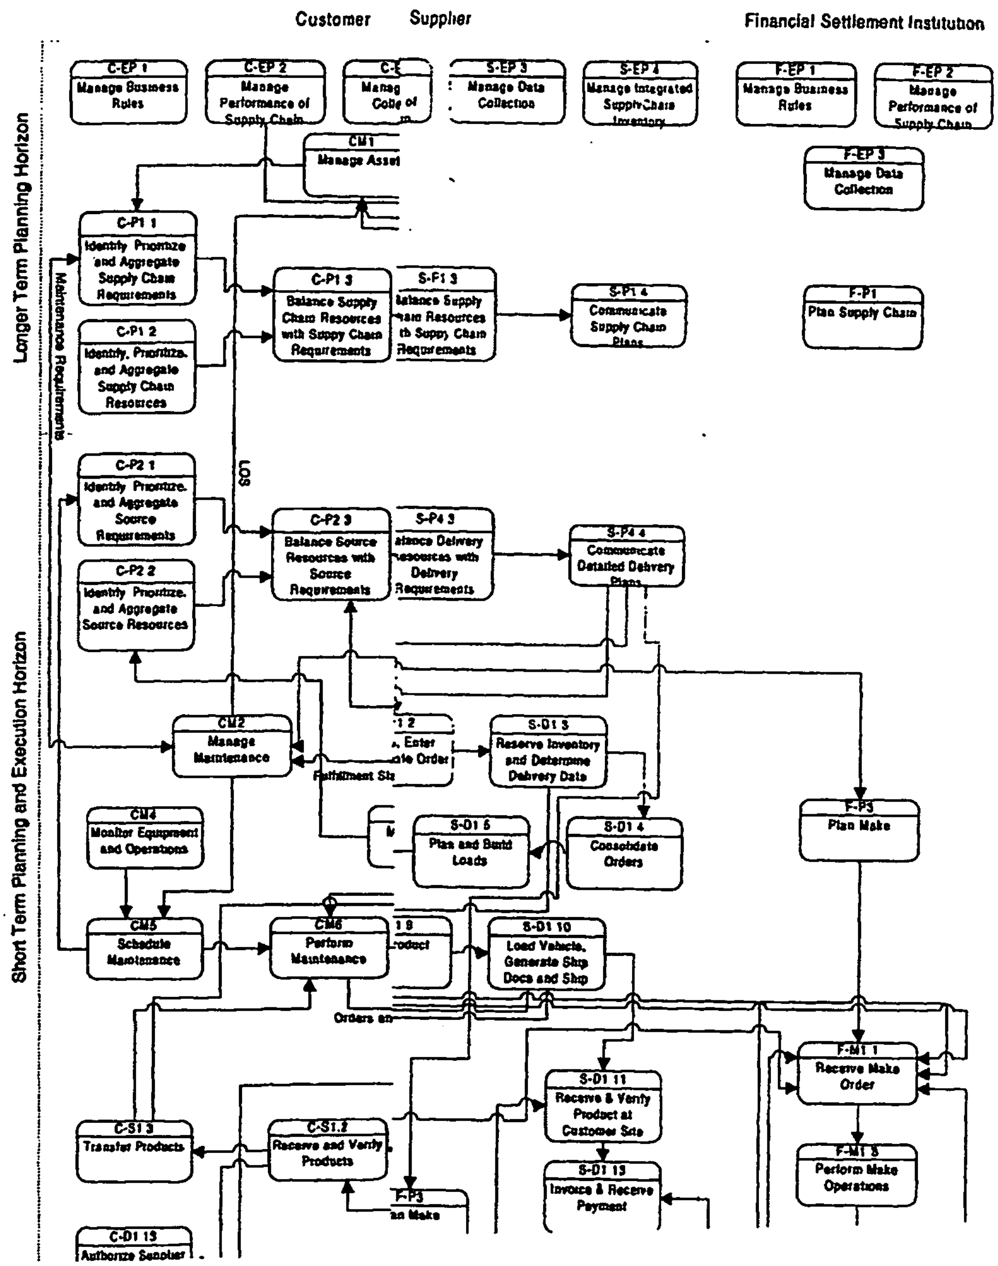

(Page 35 of 35)

No.______

Witnessed & Understood by me,	Date	Invented by	Date
		Recorded by	

CLAIMED SUBJECT MATTER

General system claims:

1. An integrated supply chain management system, comprising:

a network of computers adapted to accept a maintenance work order;

a database associated with the network of computers having stored thereon forecast data; and

programming resident on the computer network for creating an advanced demand notice comprising data representative of a product needed to perform the maintenance specified in the maintenance work order, for using the forecast data to determine the probability that the product will be needed to perform the maintenance specified in the maintenance work order, and for using the probability to select a level of fulfillment for the product.

Programming resident on the computer network for developing and executing fulfillment plans for the advance demand notices.

Additional claims directed to:

- The programming being software agents.
- The location of the software agents within the network.
- The types of software agents (e.g., MRO agents, Maintenance agents, Demand Fulfillment monitoring agents, PO monitoring agents, Supply chain performance

monitoring agents, system interface agents, fulfillment planning and excution agents, etc.).

- The types of data that may be included in forecast data (e.g., expected consumption rate data, deterministic demand data, non-deterministic demand data, historical consumption data, cause/effect data which indicates potential causes resulting in consumption).

General database claims:

1. A computer-readable media having stored thereon a data structure, comprising:

a first set of data fields containing data representing expected consumption rates for products;

a second set of data fields containing data representing deterministic demand rates for products; and

a third set of data fields containing data representing non-deterministic demand rates for products.

All three sets of data are associated with a catalog of equipment and the products needed to maintain the equipment.

General method claims:

1. A method for managing a supply chain, the method comprising:

extracting from a maintenance work order information pertaining to one or more products;

determining the probability that the one or more products will be needed to perform work in accordance with the maintenance work order; and

using the probability to determine a level of fulfillment for the one or more products.

Additional method claims directed to:

- Using a software agent to extract the information from the maintenance work order.
- Using a software agent to create an advanced demand notice.
- Using software agents to determine the probability.
- The steps and information used to determine the probability.
- The steps and information used to determine a level of fulfillment.
- The steps and information used to determine if a product is to be shipped.
- The steps and information used to determine if requisition threshold policies are being met.
- The steps and information used to determine sourcing alternatives.
- The steps and information used to determine if off-line sourcing is required.
- The steps and information used to convert the advanced demand notice to a purchase order.
- The steps and information used to monitor the status of the fulfillment progress.
- The steps and information used to monitor the process for errors and/or modifications.
- The steps and information used to monitor shipping status.
- The steps and information used to facilitate billing/invoices.

- The steps for updating the database.

Computer readable media claims:

1. A computer-readable media having instructions for managing a supply chain, the instructions performing steps comprising:

extracting from a maintenance work order information pertaining to one or more products;

determining the probability that each of the one or more products will be needed to perform work in accordance with the maintenance work order; and

using the probability to determine a level of fulfillment for the one or more products.

Intelligent Agent Architecture

Intelligent Agent and Transaction Network

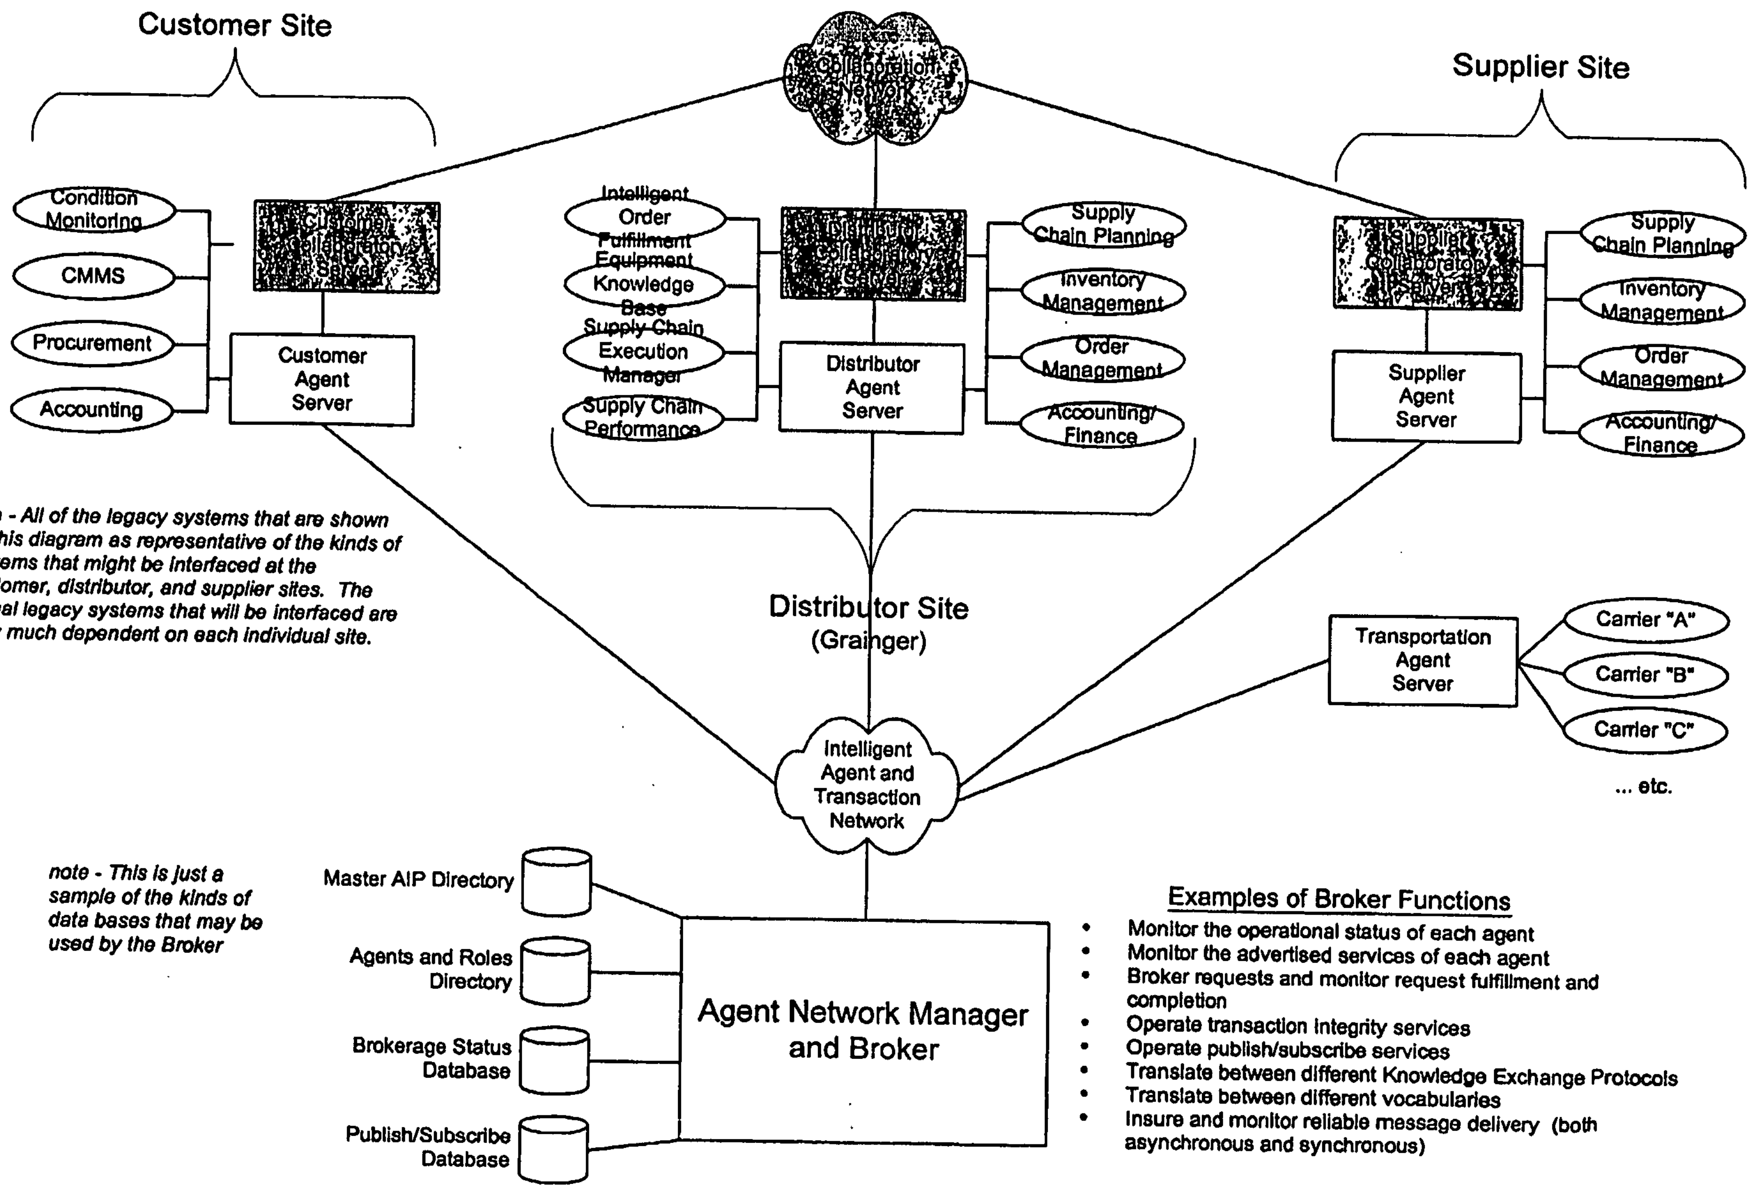

note - All of the legacy systems that are shown on this diagram as representative of the kinds of systems that might be interfaced at the customer, distributor, and supplier sites. The actual legacy systems that will be interfaced are very much dependent on each individual site.

note - This is just a sample of the kinds of data bases that may be used by the Broker

Examples of Broker Functions

- Monitor the operational status of each agent
- Monitor the advertised services of each agent
- Broker requests and monitor request fulfillment and completion
- Operate transaction integrity services
- Operate publish/subscribe services
- Translate between different Knowledge Exchange Protocols
- Translate between different vocabularies
- Insure and monitor reliable message delivery (both asynchronous and synchronous)

Diagram 1

Intelligent Agent Architecture

Agent Server Components

Depicted in the Context of a Customer Site

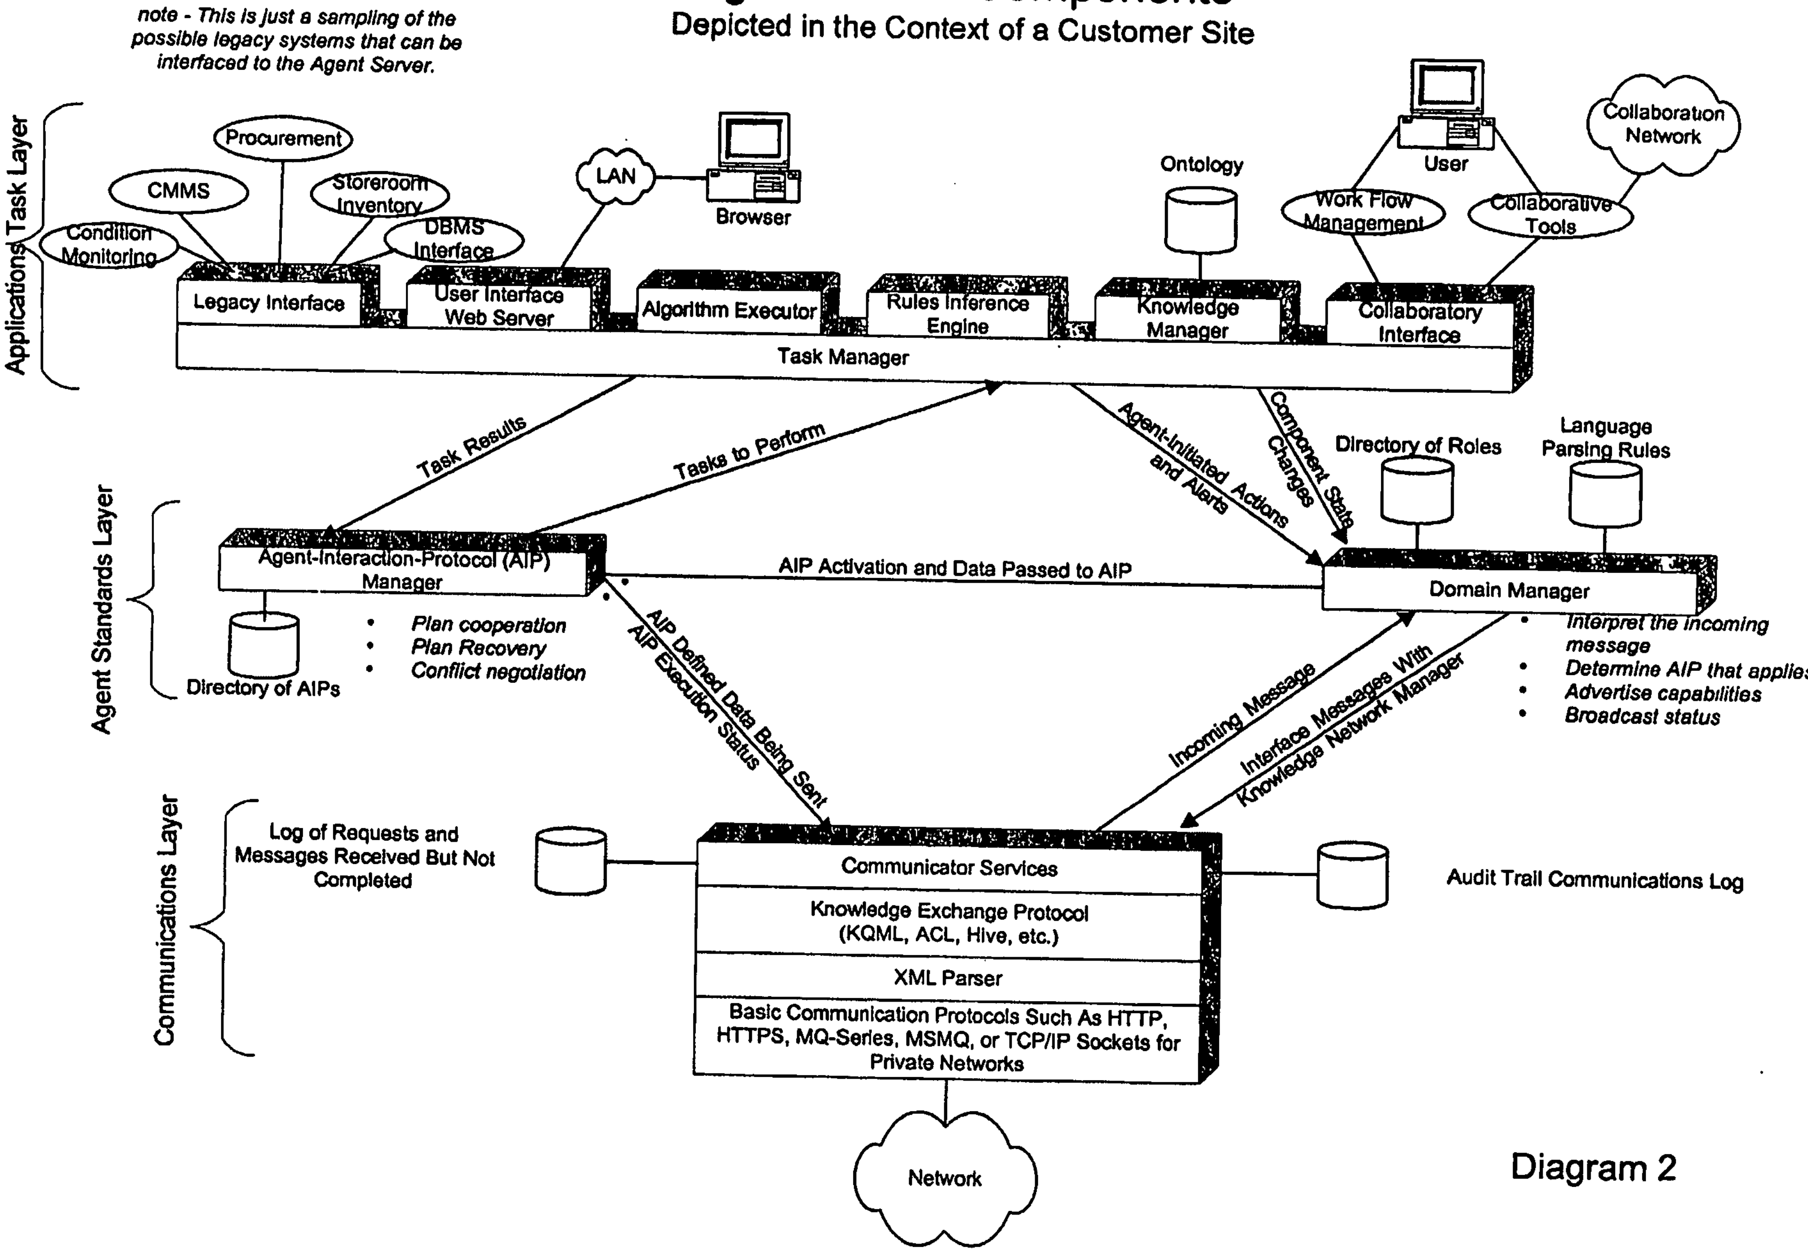

Diagram 2

Ontology Structure and Meta-Knowledge

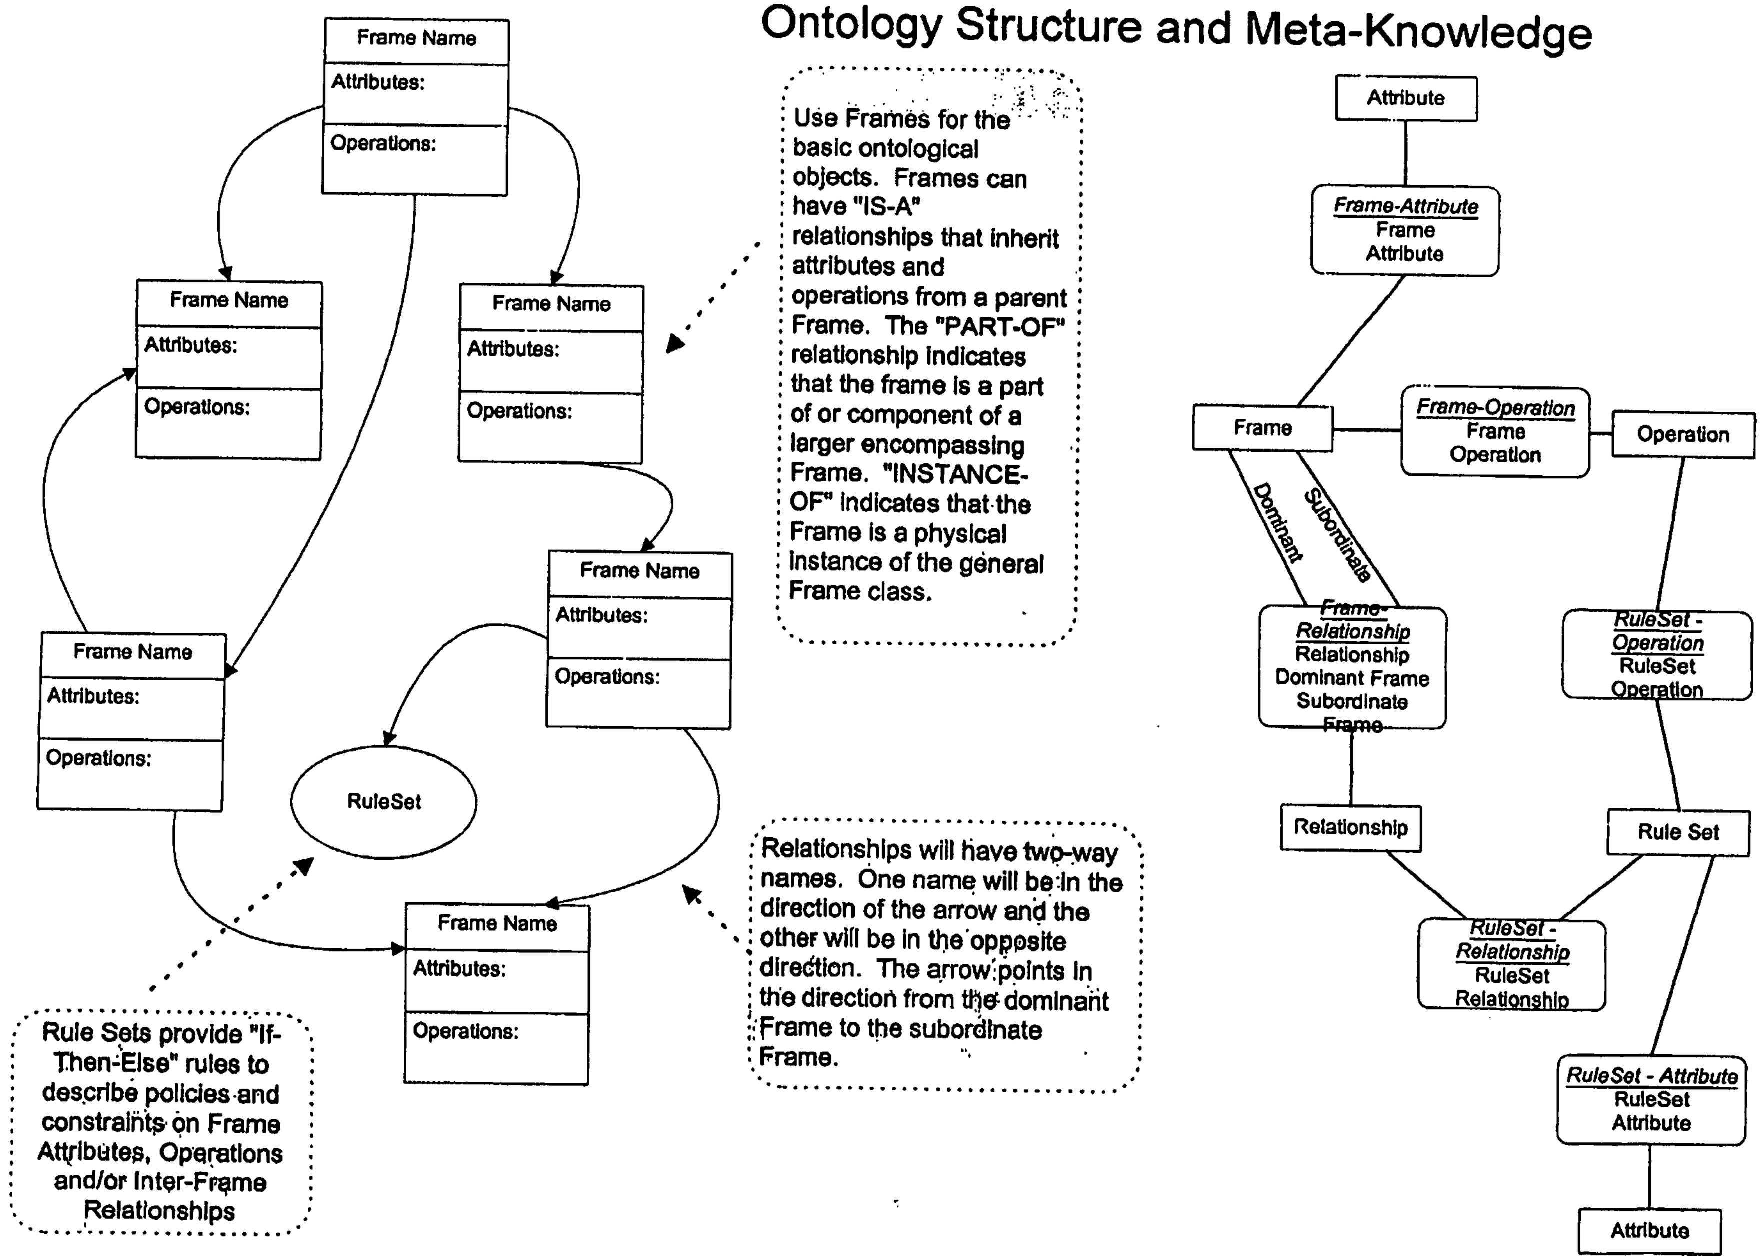

Diagram 3

Supply Chain Collaboration Points

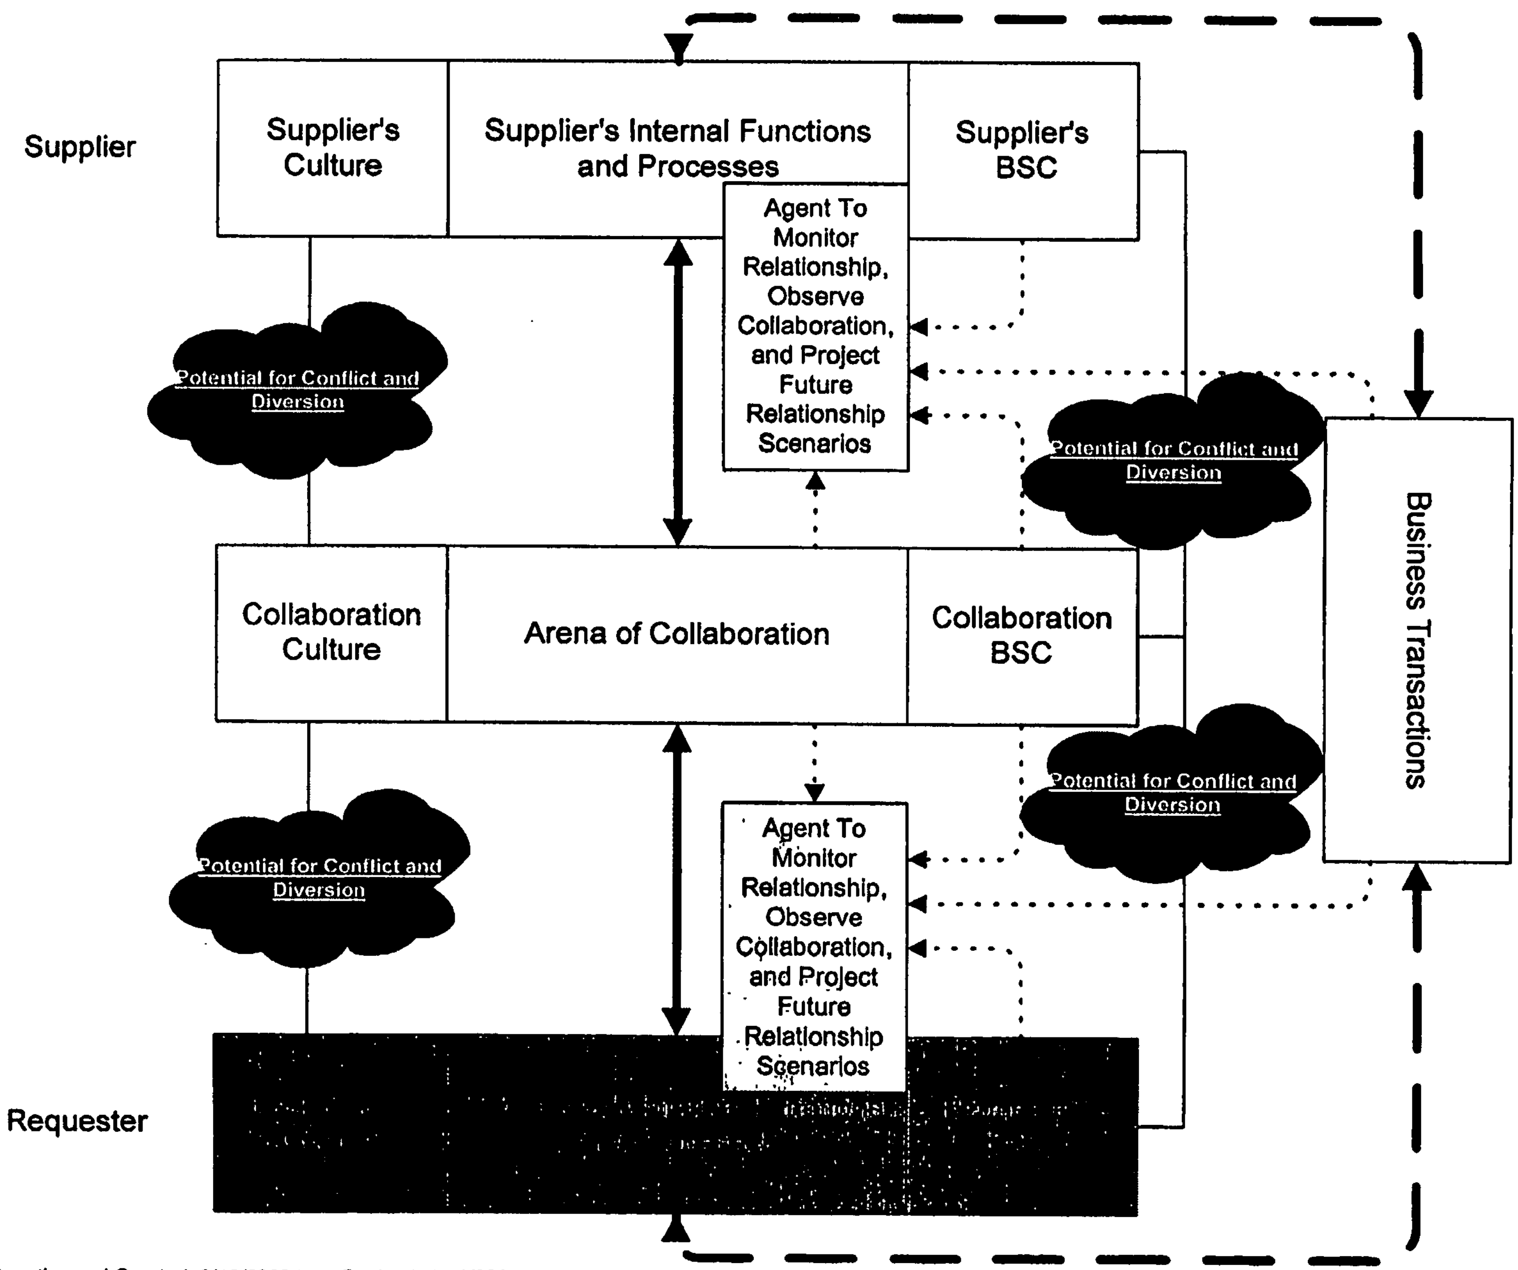

Collaboration.vsd Created: 9/13/2000 Last Revised: 11/6/2000 by Robert H. Scheer

Supply Chain Collaboration Life Cycle

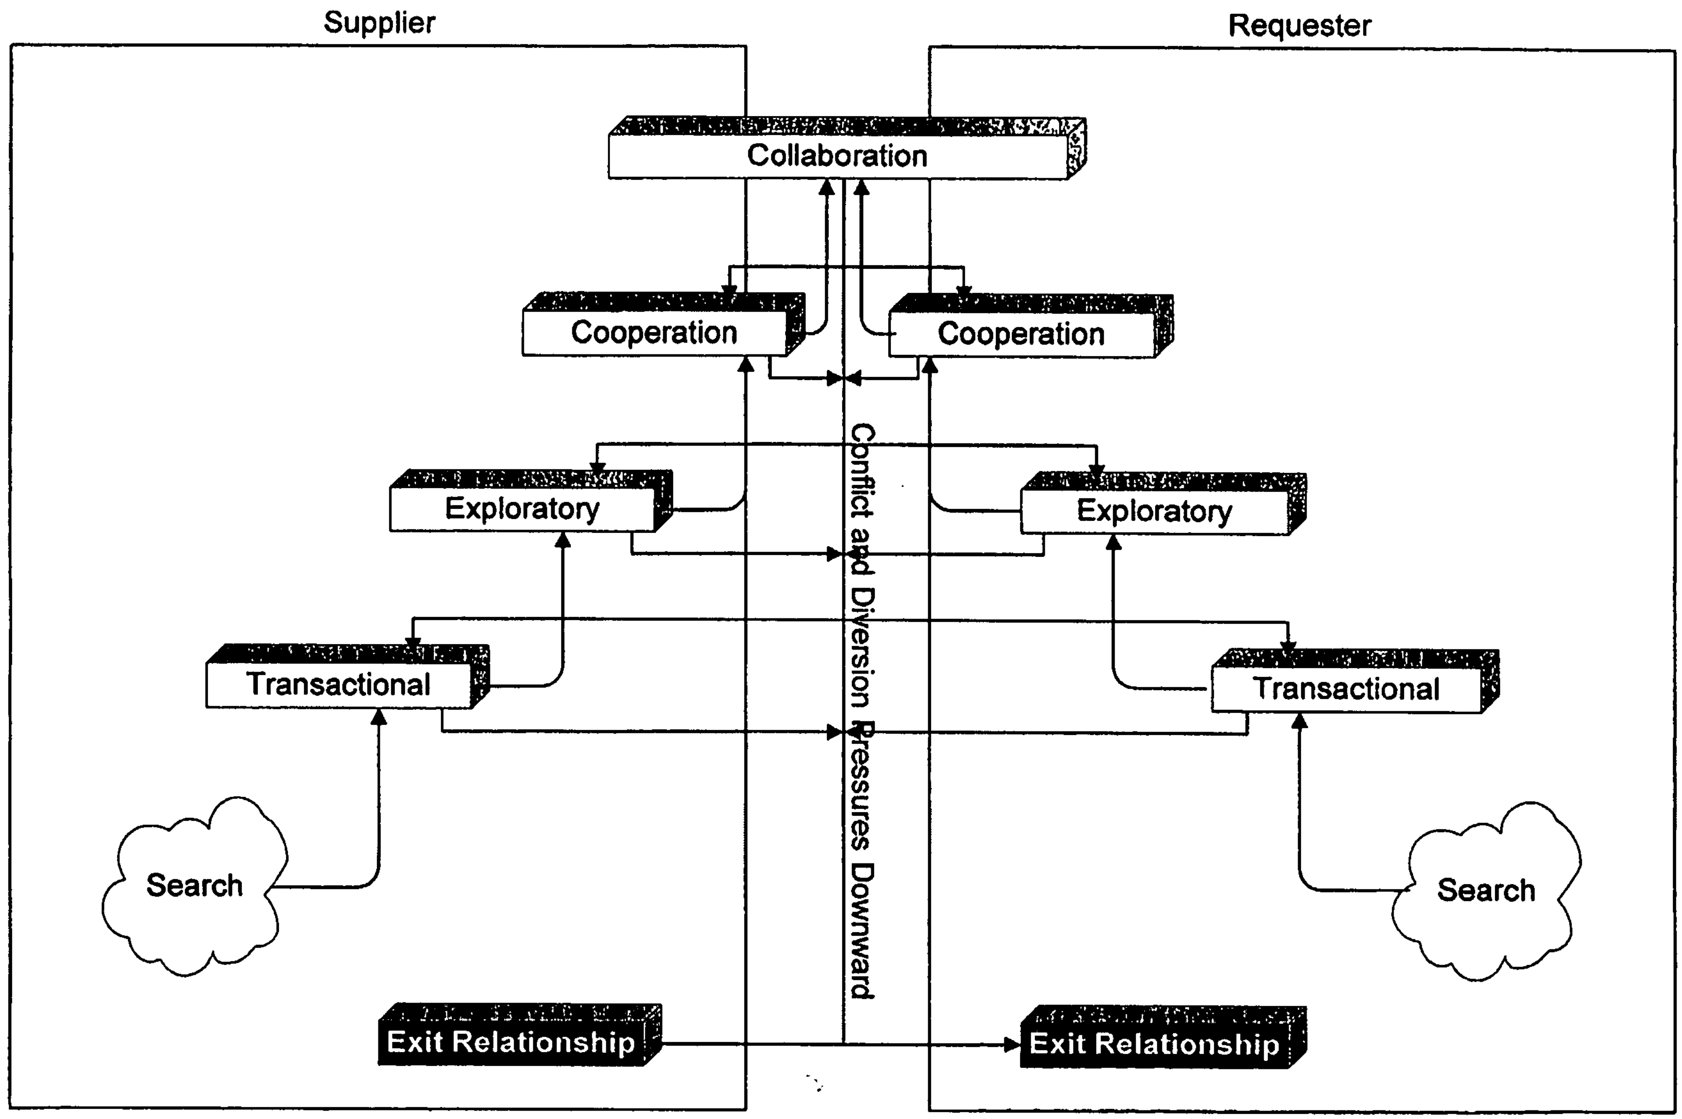

Collaboration.vsd Created: 9/13/2000 Last Revised: 11/6/2000 by Robert H. Scheer

The Collaboratory in SCOR Model Framework

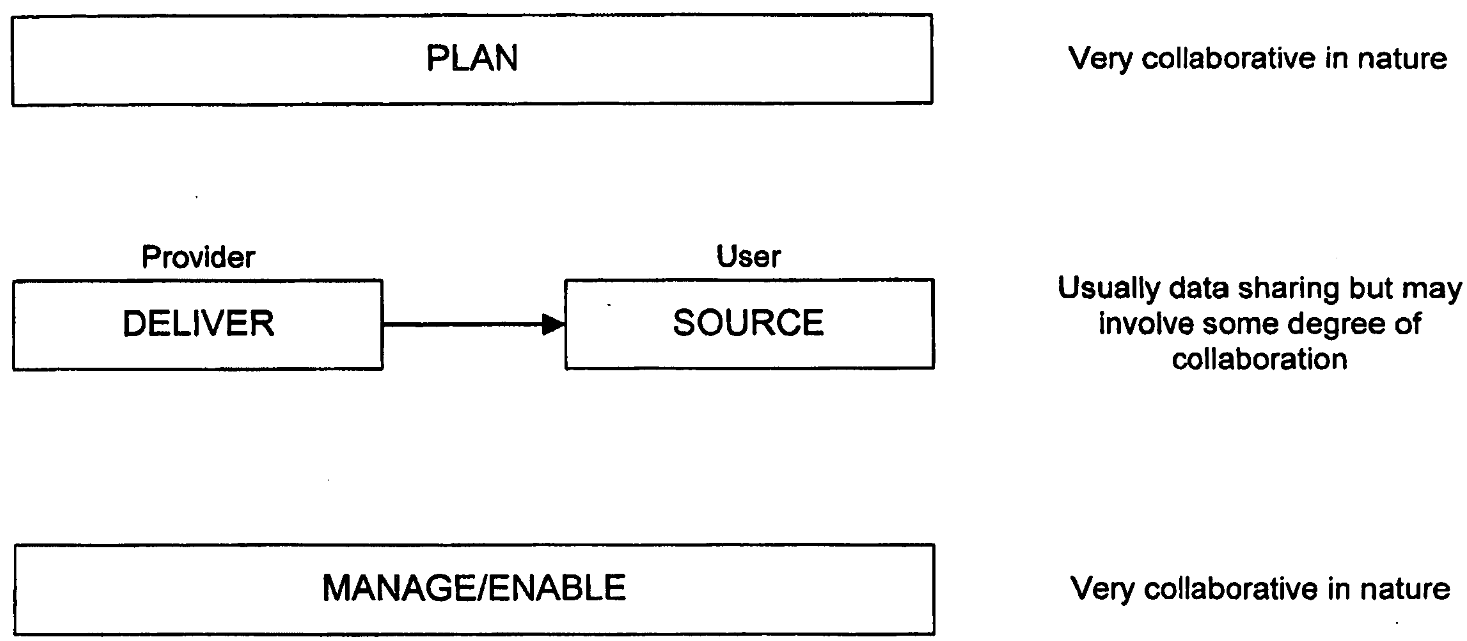

The Collaboratory Architecure

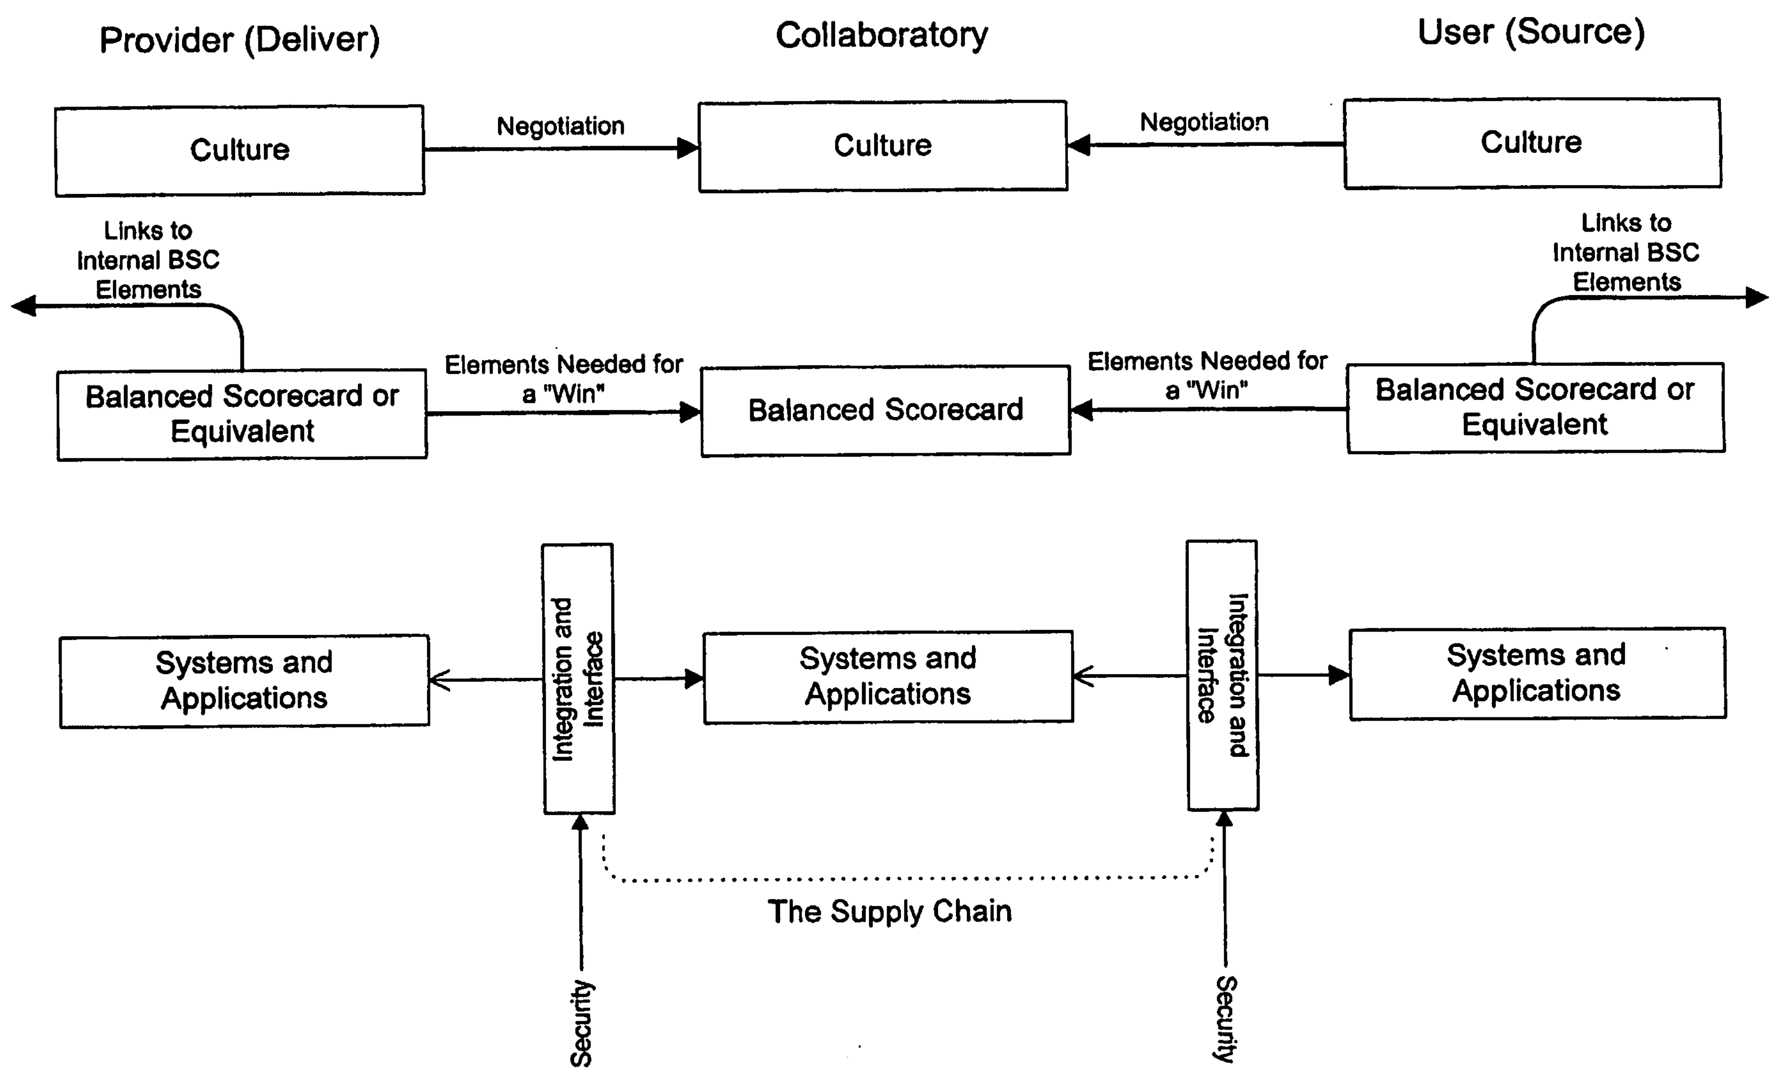

Collaboration.vsd Created: 9/13/2000 Last Revised: 11/6/2000 by Robert H. Scheer

Range of Interfirm Links

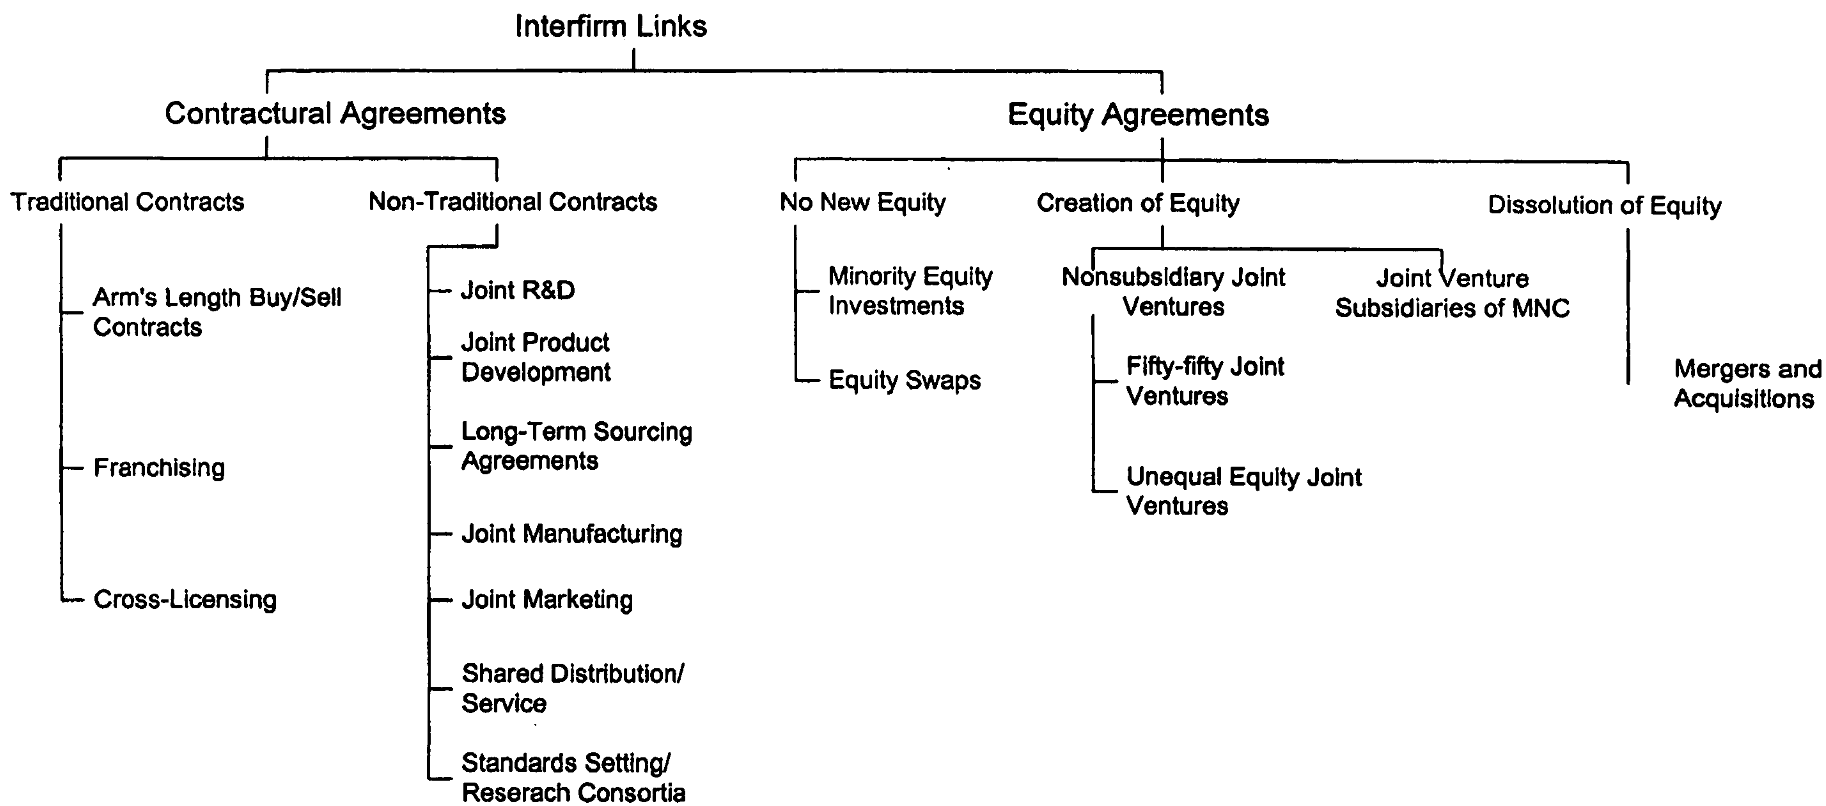

Yoshino, M. and S. Rangan (1995), "Strategic Alliances: An Entrepreneurial Approach to Globalization", Boston, MA: Harvard Business School Press (page 8)

SCOR Model Rendering

MRO Integrated Supply Chain Management -- Feature Hierarchy

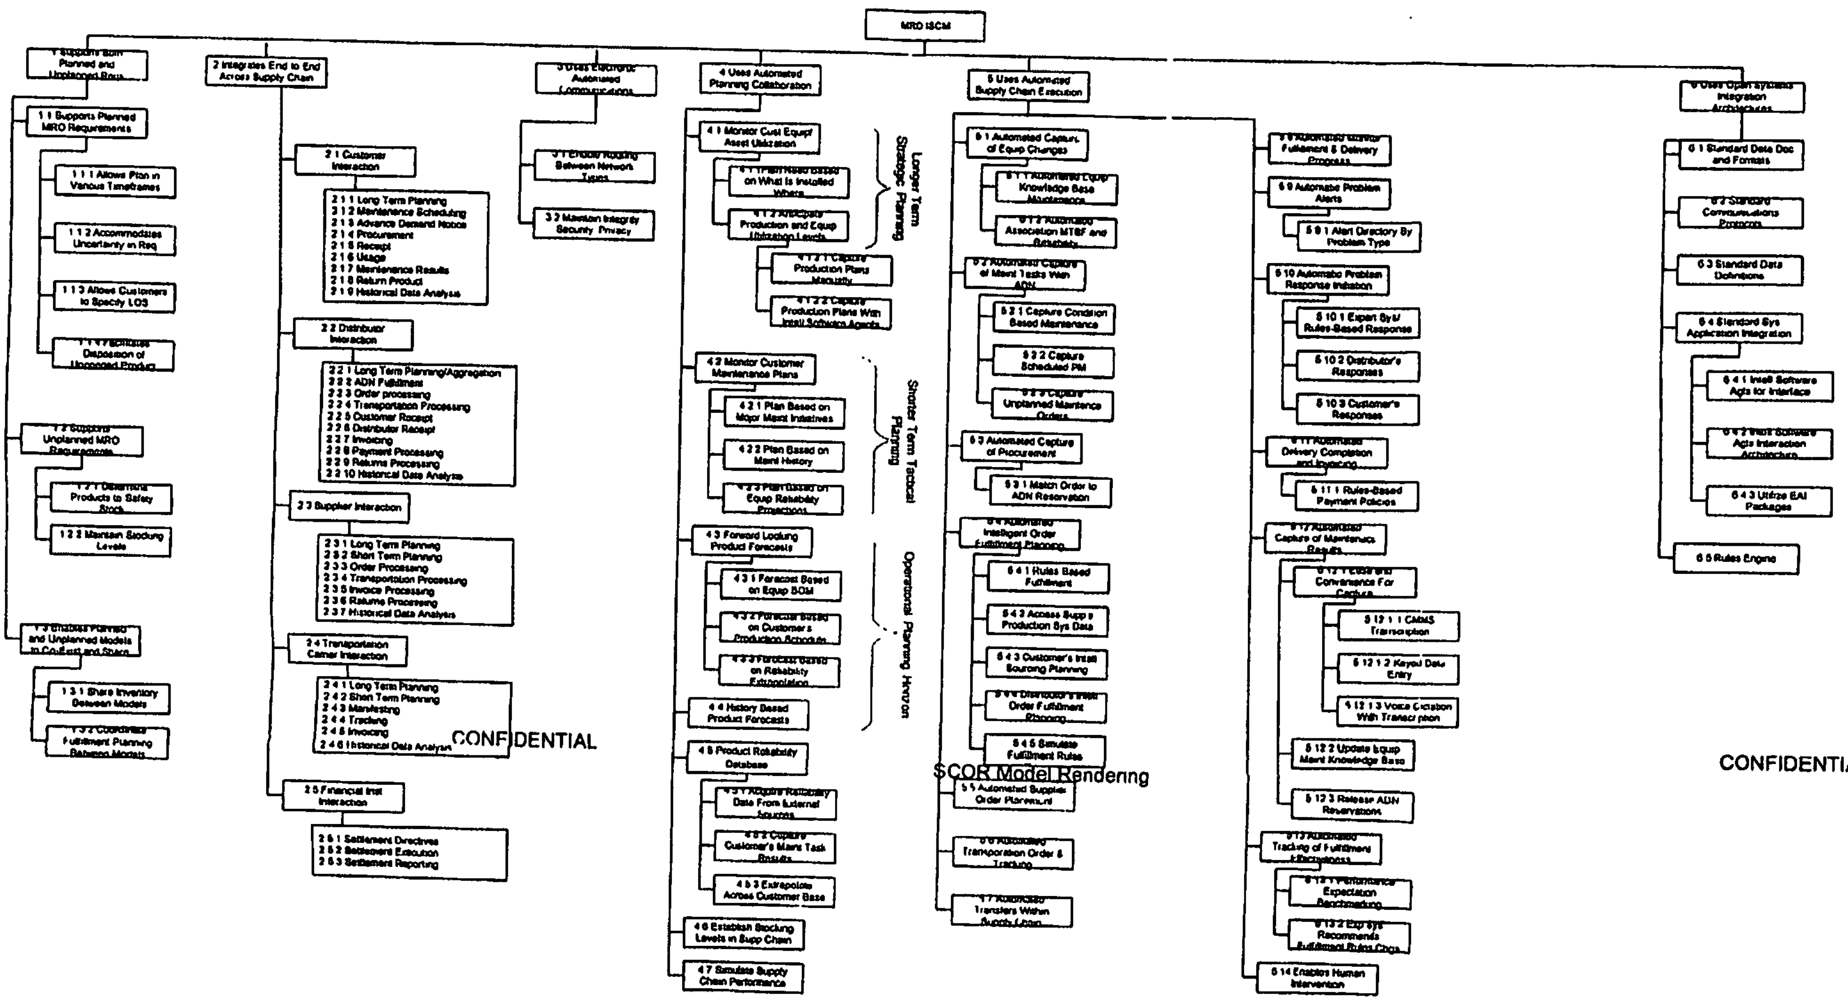

MROISCM_SCOR v40 FeatureHierarchy [rev 9/11/2000]

SCOR Model Rendering

MRO Integrated Supply Chain Management Environment and System

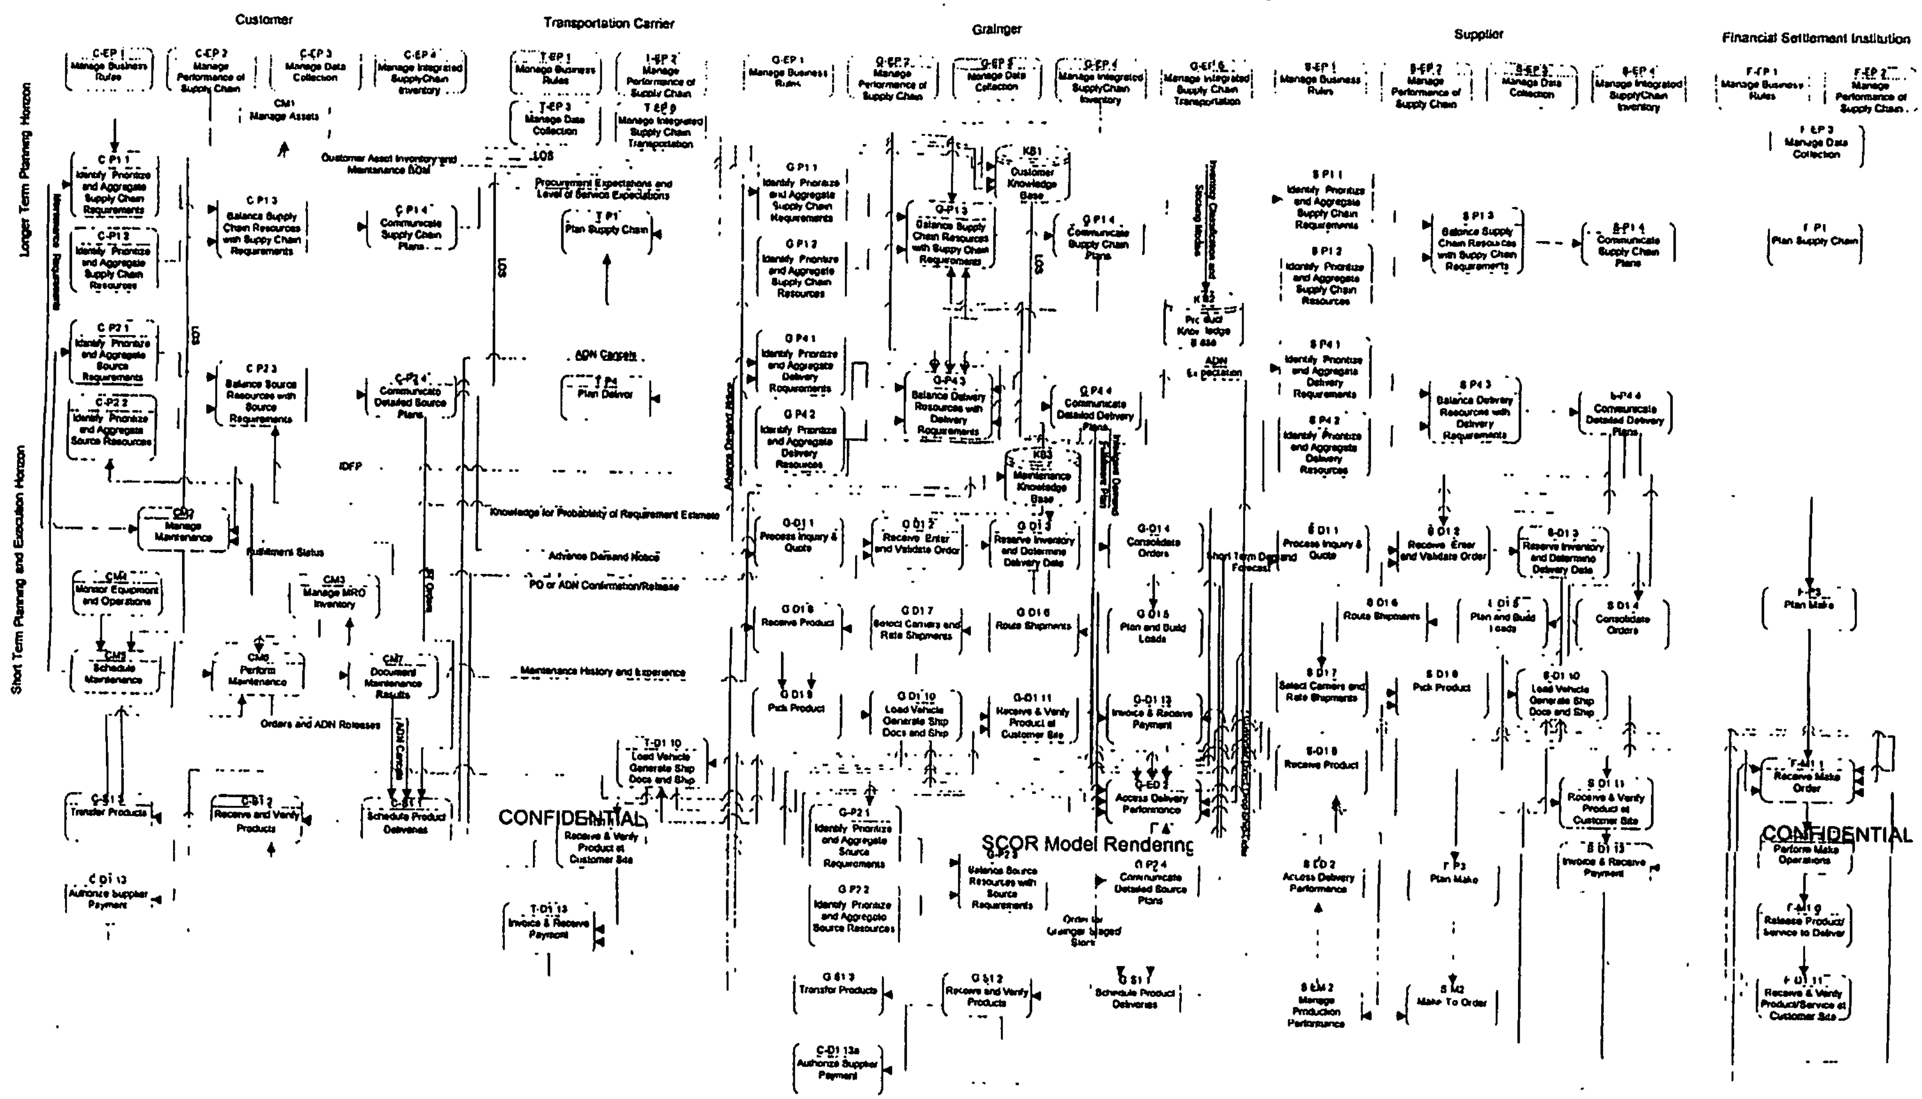

MROISCM_SCOR vsd TopLevel (rev 9/11/2000)

MRO ISCM - Equipment Knowledge Base

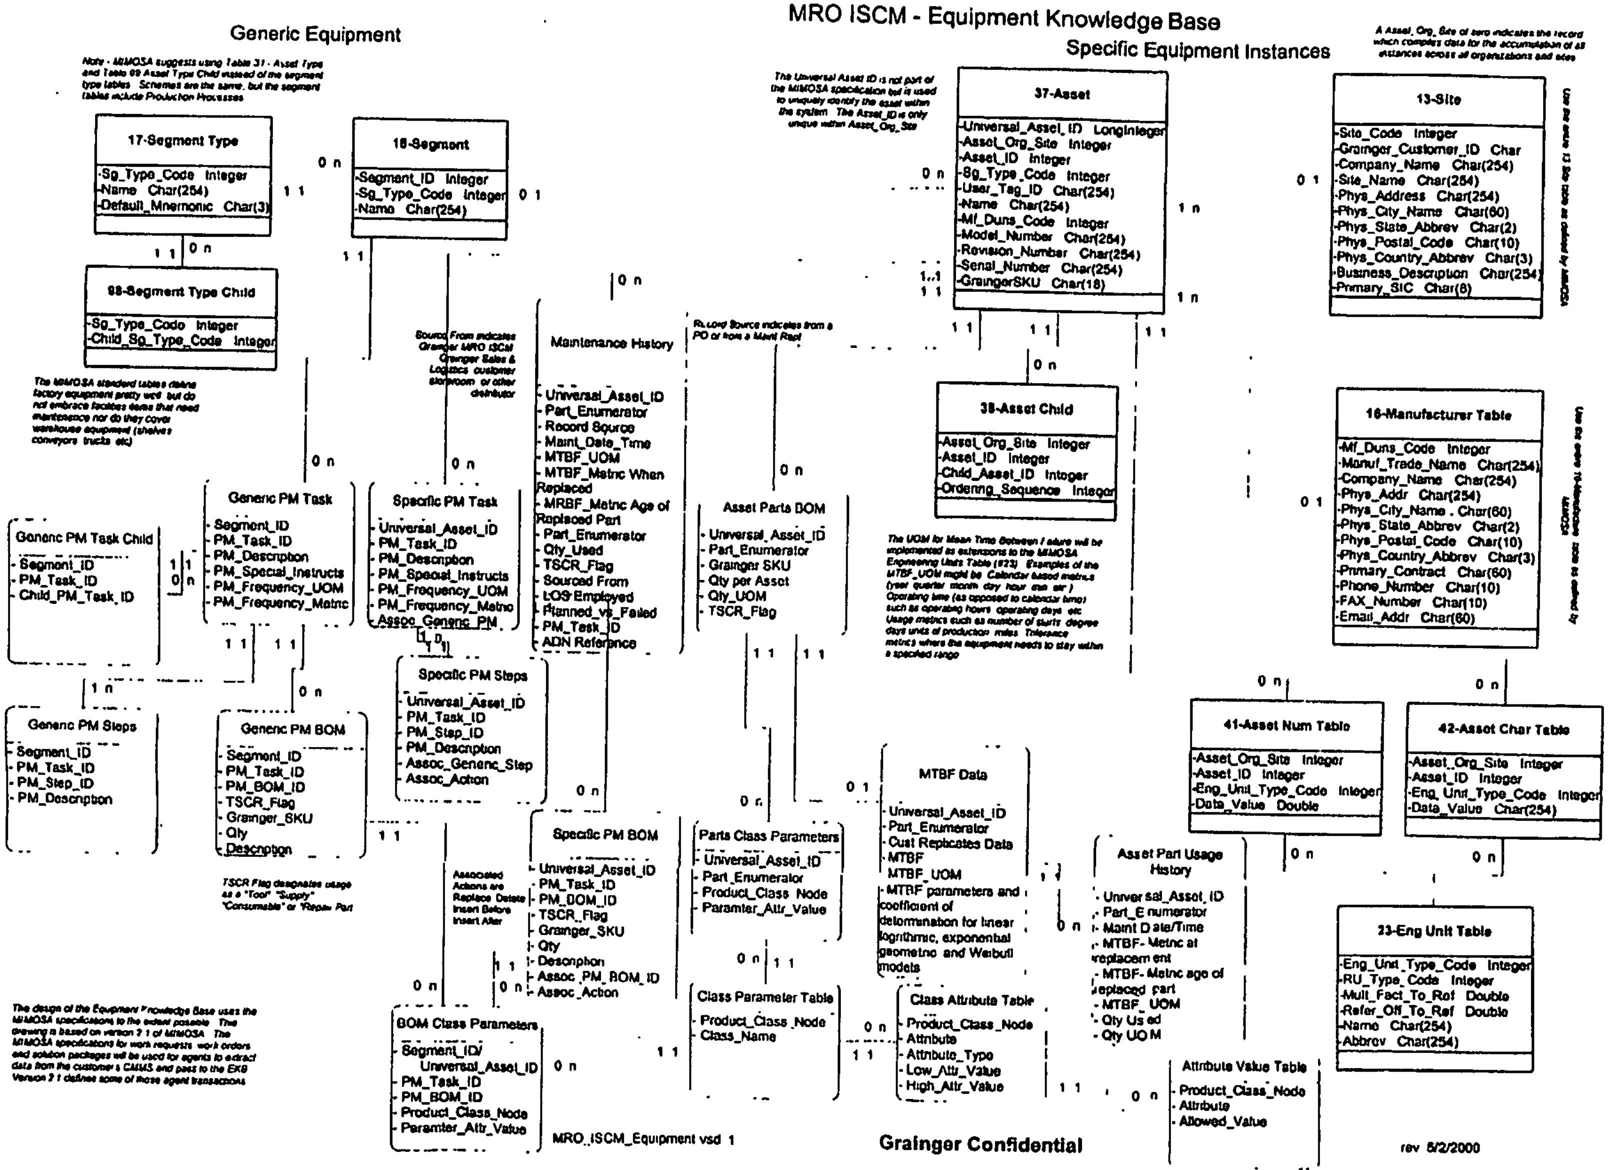

MRO_ISCM_Equipment vsd 1

Grainger Confidential

rev 5/2/2000

New Customer Equipment Data Added to Equipment Knowledge Base (cont)

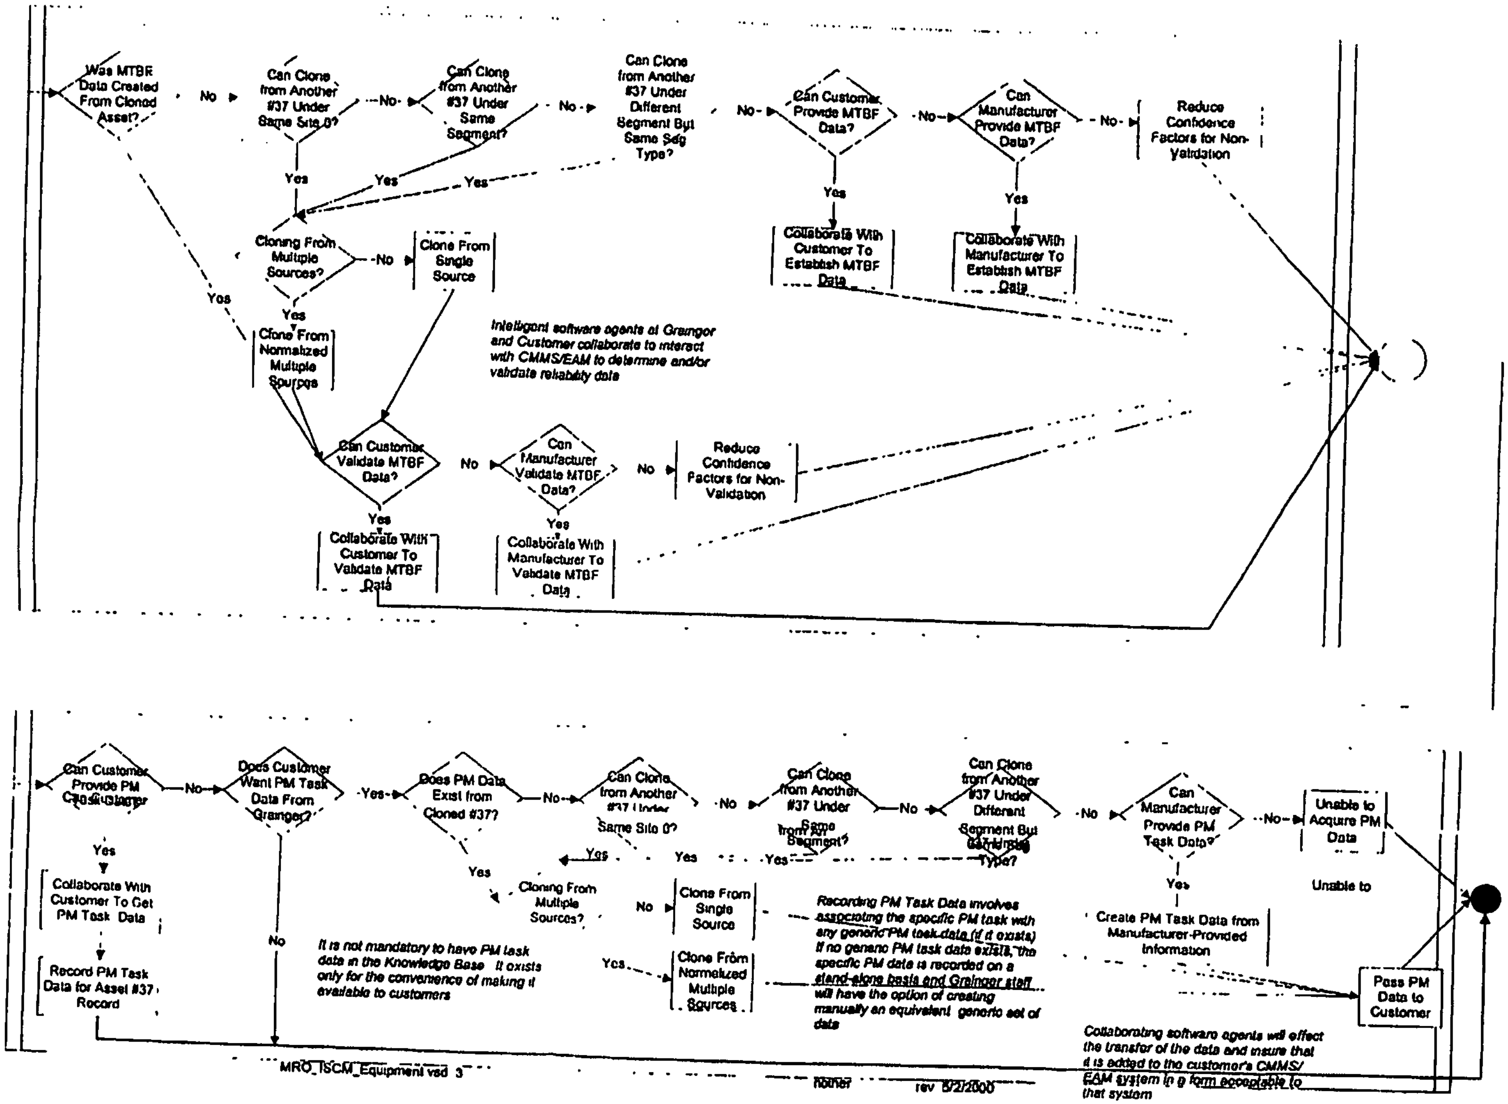

Recording PM Task Data involves associating the specific PM task with any generic PM task data (if it exists). If no generic PM task data exists, the specific PM data is recorded on a stand-alone basis and Grainger staff will have the option of creating manually an equivalent generic set of data

Iterate for Each Asset

Record Maintenance Results in the Equipment Knowledge Base

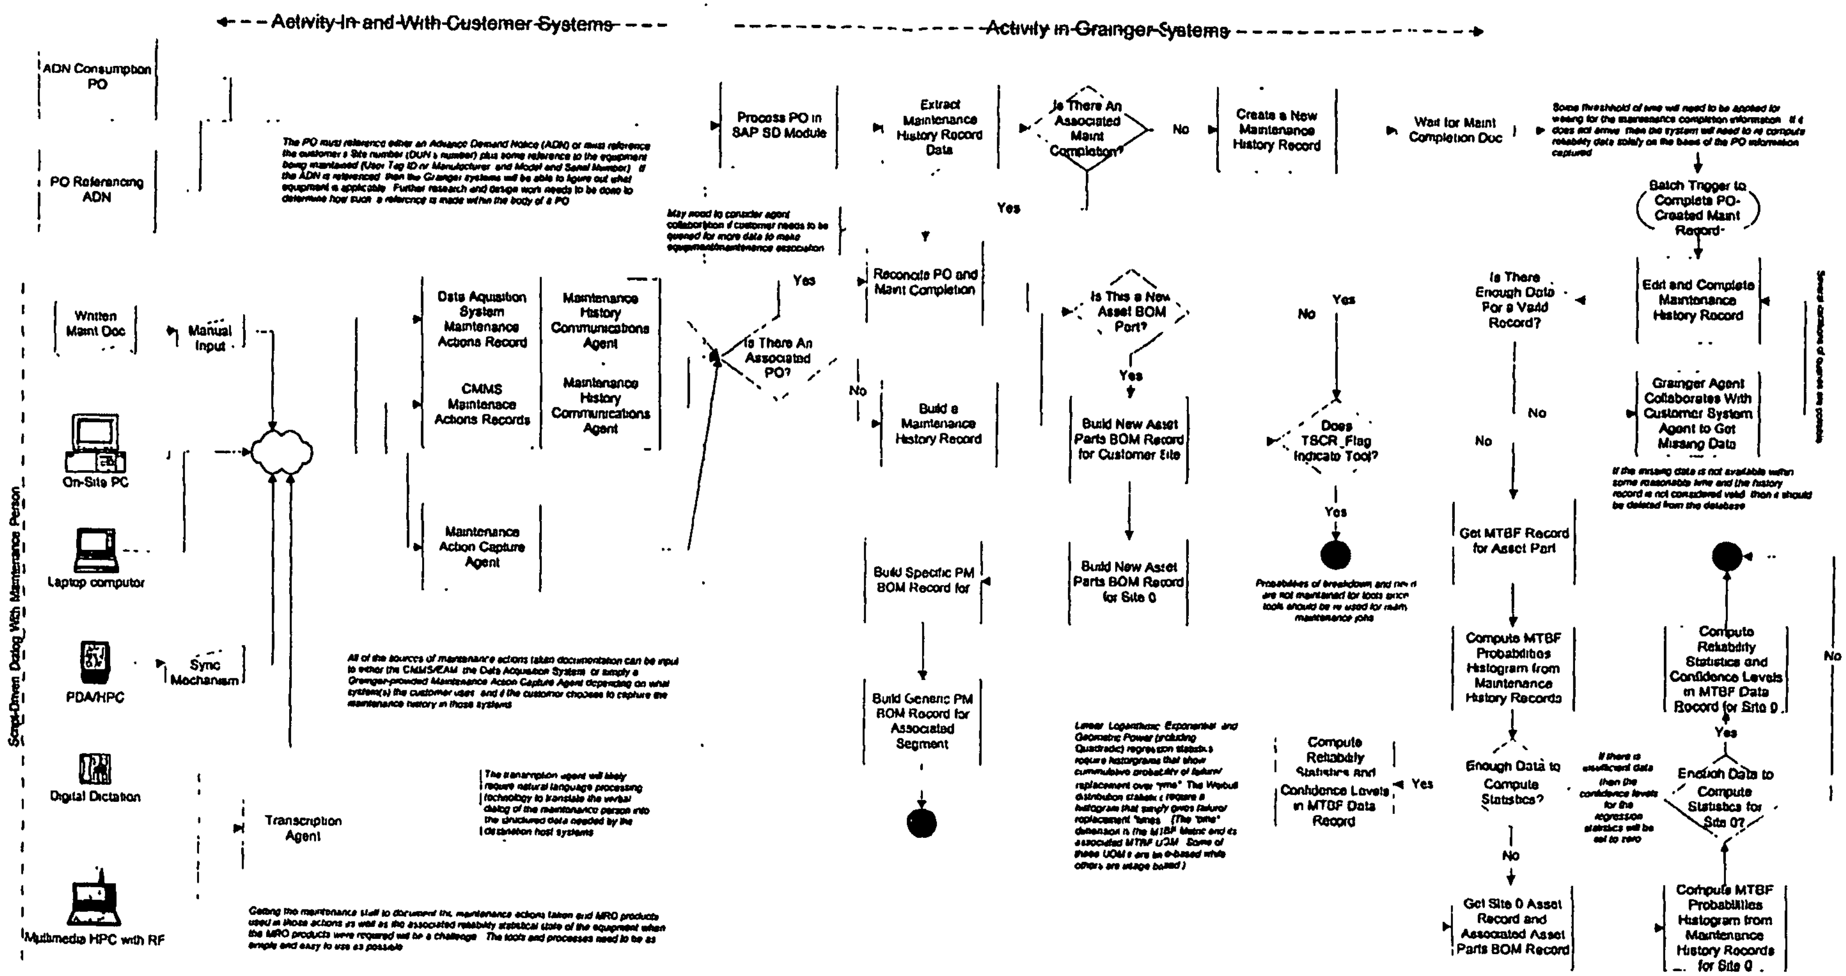

Repair Parts Reliability Computations

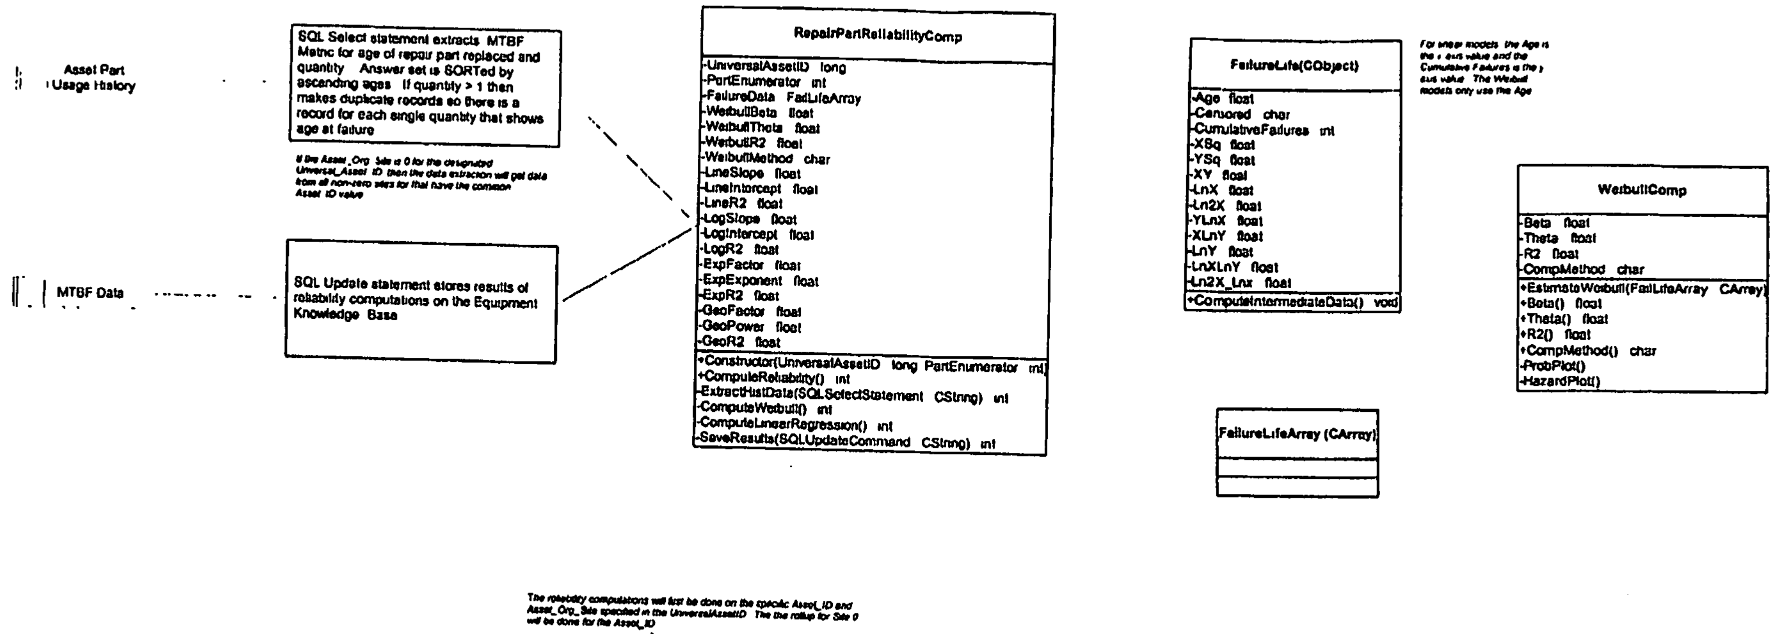

New Customer Equipment Data Added to Equipment Knowledge Base